Red Hat Enterprise Linux 10

10.2 Release Notes

Release Notes for Red Hat Enterprise Linux 10.2

Last Updated: 2026-05-20

Red Hat Enterprise Linux 10 10.2 Release Notes

Release Notes for Red Hat Enterprise Linux 10.2

Legal Notice

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 10.2 and document known issues in this release, as well as notable fixed issues, Technology Previews, deprecated functionalities, functionalities removed in RHEL 10, and other details. For information about installing Red Hat Enterprise Linux, see Installation.

Table of Contents

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We are committed to providing high-quality documentation and value your feedback. To help us improve, you can submit suggestions or report errors through the Red Hat Jira tracking system.

Procedure

1. Log in to the Jira website.
   If you do not have an account, select the option to create one.
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 OF RED HAT ENTERPRISE LINUX 10.2

Review major changes to core components and supported in-place upgrade paths in Red Hat Enterprise Linux 10.2.

1.1. MAJOR CHANGES IN RHEL 10.2

1.1.1. Security

Review the most notable changes to security in Red Hat Enterprise Linux 10.2.

The **keylime-agent** package is rebased to upstream version 0.2.9, which includes a new agent-driven push attestation model, expanded hardware cryptography, flexible TPM RSA support, and the use of ECC-signed TLS certificates.

The **clevis-pin-trustee** package provides a new Clevis pin trustee that enables automated encryption and decryption of LUKS-encrypted volumes by using remote attestation through the Trustee Key Broker Service (KBS).

The **fapolicyd** packages are rebased to upstream version 1.4.3, and you can now filter rules.

RHEL 10.2 introduces the **capnproto** package, which provides a high-performance data interchange and remote procedure call (RPC) system that uses zero-copy serialization to eliminate the overhead of traditional data encoding and decoding.

This release of the **openssh** packages introduces support for the ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism) post-quantum (PQ) key exchange combined with elliptic curves standardized by the National Institute of Standards and Technology (NIST) in FIPS mode.

Also, the **libssh** library introduces support for post-quantum traditional (PQ/T) hybrid key exchange methods based on the quantum-resistant ML-KEM standard and traditional Elliptic-curve Diffie-Hellman (ECDH) key exchange schemes.

The **p11-kit** packages have been upgraded to upstream version 0.26.1, which delivers support for post-quantum cryptography (PQC) definitions in PKCS #11 headers.

The **podman-sequoia** library supports composite post-quantum signatures.

See New features and enhancements - Security for more information.

1.1.2. Infrastructure services

Review the most notable changes to infrastructure services in Red Hat Enterprise Linux 10.2.

- PostgreSQL 18 packages are available.
- MariaDB 11.8 packages are available.
- PHP is available in version 8.4.
- The **chrony** packages are updated to version 4.8.
- The **frr** packages are updated to version 10.4.1.

For more information, see New features and enhancements - Infrastructure services .

1.1.3. Kernel

Review the most notable kernel updates in Red Hat Enterprise Linux 10.2.

- Support for Kernel Livepatches is now available in RHEL 10.
- Extends kernel observability with additional **perf** features and new Intel core, uncore, c-state, and package performance events.
- Aligns **perf** and BPF tooling more closely with upstream by updating **perf** to recent upstream versions and enabling debuginfod support.
- Expands **uncore** and **core** performance counters for newer Intel platforms and adds AMD IBS load-latency filtering to improve CPU and memory analysis.
- Adds or updates drivers and device IDs for Intel EDAC, Intel QAT, and Intel/AMD accelerator and crypto devices to improve hardware coverage.
- Improves real-time analysis and tuning by extending **rtla** threshold-overflow actions, adding **cpupower** Python bindings, and updating **rteval**.
- Updates kernel debugging and crash analysis by rebasing **crash** and enhancing LUKS-aware kdump handling in both the kernel and kdump utilities.

1.1.4. Dynamic programming languages, web and database servers

Review the most notable changes to dynamic programming languages, web and database servers in Red Hat Enterprise Linux 10.2.

Later versions of the following Application Streams are now available:

- **Node.js 24**

Later versions of the following web servers are now available:

- **Apache HTTP Server 2.4.63**

Later versions of the following database servers are now available:

- **MariaDB 11.8**

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

1.1.5. Compilers and development tools

Review the most notable changes to compilers and development tools in Red Hat Enterprise Linux 10.2.

System toolchain

The following system toolchain components are available with RHEL 10.2:

- **GCC 14.3**
- **glibc 2.39**
- **Annobin 13.02**

- **Binutils 2.41**

Performance tools and debuggers

The following performance tools and debuggers are available with RHEL 10.2:

- **GDB 16.3**
- **Valgrind 3.26.0**
- **SystemTap 5.4**
- **Dyninst 13.0.0**
- **elfutils 0.194**
- **libabigail 2.9**

Performance monitoring tools

The following performance monitoring tools are available with RHEL 10.2:

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

Compiler toolsets

The following compiler toolsets are available with RHEL 10.2:

- **GCC Toolset 15**
  - **GCC 15.2**
  - **Binutils 2.44**
    Note that **Annobin** and **dwz** are not provided in GCC Toolset starting with version 15.
- **LLVM Toolset 21.1.8**
- **Rust Toolset 1.92.0**
- **Go Toolset 1.26.2**

1.1.6. Desktop

Review the most notable changes to desktop in Red Hat Enterprise Linux 10.2.

Flatpaks are the default delivery method for Mozilla Firefox and Thunderbird. The default delivery method for Mozilla Firefox and Thunderbird is changed from RPM packages to Flatpaks. Anaconda, the RHEL installer, preinstalls these Flatpaks by default.

See New features and enhancements - Desktop for more information.

1.2. IN-PLACE UPGRADE

Review the most notable changes to in-place upgrades in Red Hat Enterprise Linux 10.2.

1.2.1. In-place upgrade from RHEL 9 to RHEL 10

The supported in-place upgrade paths currently are:

- From RHEL 9.6 to RHEL 10.0 and 9.8 to 10.2 on the following architectures:
  - AMD and Intel 64-bit architectures (x86-64-v3)
  - The 64-bit ARM architecture (ARMv8.0-A)
  - IBM Power Systems, Little Endian (POWER10) and later
  - 64-bit IBM Z (IBM z15 or IBM LinuxONE III or later)

For instructions on performing an in-place upgrade, see Upgrading from RHEL 9 to RHEL 10 .

For information regarding how Red Hat supports the in-place upgrade process, see the In-place upgrade Support Policy.

Notable enhancements and bug fixes include:

- New Ansible system roles to automate the upgrade process. For more information, see In-place upgrade phases automation with the **analysis**, **remediate**, and **upgrade** Ansible roles.
- Modernization of the system storage initialization when booting to the upgrade environment.
- Correctly upgrade systems with configured LVM and deliver initial improvements for multipath.
- Fix the upgrade on systems with Non-Volatile Memory Express over Fibre Channel (NVMe-FC).
- Preserve Network Interface Card (NIC) names during the upgrade by using the net.naming-scheme argument in the kernel command line.
- Migrate kerberos configuration during the upgrade.

1.2.2. In-place upgrade from RHEL 8 to RHEL 10

It is not possible to perform an in-place upgrade directly from RHEL 8 to RHEL 10. However, you can perform an in-place upgrade from RHEL 8 to RHEL 9 and then perform a second in-place upgrade to RHEL 10. For more information, see In-place upgrades over multiple RHEL major versions by using Leapp.

1.3. RED HAT CUSTOMER PORTAL LABS

Review the most popular Red Hat Customer Portal Labs in Red Hat Enterprise Linux 10.2.

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

- Registration Assistant
- Kickstart Generator
- Red Hat Product Certificates

- Red Hat CVE Checker
- Kernel Oops Analyzer
- Red Hat Satellite Upgrade Helper
- Load Balancer Configuration Tool
- Ceph Placement Groups (PGs) per Pool Calculator
- Red Hat Out of Memory Analyzer
- Postfix Configuration Helper
- Red Hat IdM Upgrade Helper
- NetworkManager Command Generator

1.4. ADDITIONAL RESOURCES

Review additional resources to effectively plan and manage your Red Hat Enterprise Linux 10 deployments. The following list includes content about system capabilities, life cycles, application compatibility, upgrade paths, troubleshooting, and other important information.

Capabilities and limits of Red Hat Enterprise Linux 10 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 10, including licenses and application compatibility levels.

Application compatibility levels are explained in the Red Hat Enterprise Linux 10: Application Compatibility Guide document.

Major **differences between RHEL 9 and RHEL 10**, including removed functionality, are documented in Considerations in adopting RHEL 10 .

Instructions on how to perform an **in-place upgrade from RHEL 9 to RHEL 10** are provided in Upgrading from RHEL 9 to RHEL 10 .

Using **Red Hat Lightspeed** you can proactively identify, examine, and resolve known technical issues. Red Hat Lightspeed is included with all RHEL subscriptions. 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 FOR RED HAT ENTERPRISE LINUX 10.2

Review the supported hardware architectures for Red Hat Enterprise Linux 10.2 to verify hardware compatibility.

Red Hat Enterprise Linux 10.2 is distributed with the kernel version 6.12.0-211.7.1, which provides support for the following architectures at the minimum required version (stated in parentheses):

- AMD and Intel 64-bit architectures (x86-64-v3)
- The 64-bit ARM architecture (ARMv8.0-A)
- IBM Power Systems, Little Endian (POWER10 and later)
- 64-bit IBM Z (z15 and later)

Make sure you purchase the appropriate subscription for each architecture.

CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 10

Review how Red Hat Enterprise Linux 10 distributes content so you can effectively plan your system deployments.

3.1. INSTALLATION

Red Hat Enterprise Linux 10 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.

3.2. REPOSITORIES

Red Hat Enterprise Linux 10 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.

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 10 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 following formats:

- RPM format
- Software Collections
- Flatpaks

NOTE

In previous RHEL major versions, some Application Streams were available as modules as an extension to the RPM format. In RHEL 10, Red Hat does not intend to provide any Application Streams that use modularity as the packaging technology and, therefore, no modular content is being distributed with RHEL 10.

Each Application Stream component has a given life cycle, either the same as RHEL 10 or shorter.

RHEL 10 improves the Application Streams experience by providing initial Application Stream versions that can be installed as RPM packages using the **dnf install** command.

NOTE

Certain initial Application Streams in the RPM format have a shorter life cycle than Red Hat Enterprise Linux 10.

Always determine what version of an Application Stream you want to install.

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.

CHAPTER 4. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS

Review major changes in the Red Hat Enterprise Linux 10.2 kernel to understand impacts on your system configuration. These changes could include, for example, updates to **proc entries**, **sysctl** and **sysfs** default values, boot parameters, and other kernel options.

4.1. NEW KERNEL PARAMETERS

microcode=

[X86] Control the behavior of the microcode loader.
Options:

- **base_rev=X**: Set the base microcode revision of each thread when in debug mode, where **X** is an unsigned 32-bit integer.
- **dis_ucode_ldr**: Disable the microcode loader.
- **force_minrev**: Enable or disable the microcode minimal revision enforcement for the runtime microcode loader.

nvme.quirks=

[NVME] Provide a list of quirk entries to augment the built-in NVMe quirk list.
Format: **nvme.quirks=<VendorID>:<ProductID>:<quirk_names>-…**

The vendor ID and product ID are 4-digit hexadecimal numbers. The **quirk_names** field is a comma-separated list of quirk names. You can prefix a quirk name with **^** to disable that quirk explicitly.

Example: **nvme.quirks=7710:2267:bogus_nid,^identify_cns-9900:7711:broken_msi**

rcutorture.gp_cond_exp=

[KNL] Use conditional or asynchronous update-side expedited-grace-period primitives, if available.

rcutorture.gp_cond_full=

[KNL] Use conditional or asynchronous update-side normal-grace-period primitives that also take concurrent expedited grace periods into account, if available.

rcutorture.gp_cond_exp_full=

[KNL] Use conditional or asynchronous update-side expedited-grace-period primitives that also take concurrent normal grace periods into account, if available.

rcutorture.gp_cond_wi=

[KNL] Set the nominal wait interval for normal conditional grace periods, in microseconds. The kernel randomly selects the actual wait interval up to this value with nanosecond granularity.
This parameter controls the wait interval for conditional grace periods that you specify through the **rcutorture.gp_cond** and **rcutorture.gp_cond_full** module parameters.

Defaults to 16 jiffies; for example, 16,000 microseconds on a system with **HZ=1000**.

rcutorture.gp_cond_wi_exp=

[KNL] Set the nominal wait interval for expedited conditional grace periods, in microseconds. The kernel randomly selects the actual wait interval up to this value with nanosecond granularity.

This parameter controls the wait interval for expedited conditional grace periods that you specify through the **rcutorture.gp_cond_exp** and **rcutorture.gp_cond_exp_full** module parameters.

Defaults to 128 microseconds.

rcutorture.gp_poll=

[KNL] Use polled update-side normal-grace-period primitives, if available.

rcutorture.gp_poll_exp=

[KNL] Use polled update-side expedited-grace-period primitives, if available.

rcutorture.gp_poll_full=

[KNL] Use polled update-side normal-grace-period primitives that also take concurrent expedited grace periods into account, if available.

rcutorture.gp_poll_exp_full=

[KNL] Use polled update-side expedited-grace-period primitives that also take concurrent normal grace periods into account, if available.

rcutorture.gp_poll_wi=

[KNL] Set the nominal wait interval for normal conditional grace periods when using polled update-side primitives, in microseconds. The kernel randomly selects the actual wait interval up to this value with nanosecond granularity.
This parameter controls the wait interval for conditional grace periods that you specify through the **rcutorture.gp_poll** and **rcutorture.gp_poll_full** module parameters.

Defaults to 16 jiffies; for example, 16,000 microseconds on a system with **HZ=1000**.

rcutorture.gp_poll_wi_exp=

[KNL] Set the nominal wait interval for expedited conditional grace periods when using polled update-side primitives, in microseconds. The kernel randomly selects the actual wait interval up to this value with nanosecond granularity.
This parameter controls the wait interval for expedited conditional grace periods that you specify through the **rcutorture.gp_poll_exp** and **rcutorture.gp_poll_exp_full** module parameters.

Defaults to 128 microseconds.

rcutorture.gpwrap_lag=

[KNL] Enable grace-period wrap lag testing in **rcutorture**. Set this parameter to false to prevent the **gpwrap** lag test from running.
The default value is **true**.

rcutorture.gpwrap_lag_gps=

[KNL] Set the number of grace periods to tolerate between the per-CPU RCU data structure (**rdp**) and the RCU node (**rnp**) **gp_seq** values before setting the overflow flag during active wrap-lag testing.
The default value is **8**.

rcutorture.gpwrap_lag_cycle_mins=

[KNL] Set the total cycle duration for **gpwrap** lag testing, in minutes. The cycle includes both active and inactive testing periods.
The default value is **30** minutes.

rcutorture.gpwrap_lag_active_mins=

[KNL] Set the duration, in minutes, during which **gpwrap** lag is active in each testing cycle. During the active period, the grace-period wrap lag is controlled by the value of **rcutorture.gpwrap_lag_gps**. The default value is **5** minutes.

rcutorture.preempt_duration=

[KNL] Set the duration, in milliseconds, of preemptions by a high-priority FIFO real-time task. Set this parameter to **0** (the default) to disable the preemption test. The kernel selects CPUs to preempt randomly from the set of CPUs that are online at that moment. Races with CPUs going offline are ignored; in such cases, that preemption attempt is skipped.

rcutorture.preempt_interval=

[KNL] Set the interval, in milliseconds, between preemptions by a high-priority FIFO real-time task. The interval defaults to 1 second.
This delay is driven by an **hrtimer** and is fuzzed to avoid unintended synchronizations between preemptions.

rcutorture.reader_flavor=

[KNL] Set a bit mask that indicates which RCU readers to use in **rcutorture**.
If more than one bit is set, **rcutorture** enters readers in order from the lowest-order bit to the highest-order bit and exits readers in the reverse order. For SRCU, the bits have the following meanings:

- **0x1** - normal readers
- **0x2** - NMI-safe readers
- **0x4** - lightweight readers

rcutorture.test_boost_holdoff=

[KNL] Set the holdoff time, in seconds, from the start of a **rcutorture** test to the start of RCU priority-boost testing.
The default value is **0**, which disables any holdoff period.

tsa=

[X86] Control mitigation for transient scheduler attacks on AMD CPUs.
For additional technical background, search for the document titled **"Technical guidance for mitigating transient scheduler attacks"**.

Values:

- **off**: Disable the mitigation.
- **on** (default): Enable the mitigation.
- **user**: Mitigate only user-to-kernel transitions.
- **vm**: Mitigate only guest-to-host transitions.

vmscape=

[X86] Control mitigation for VMscape attacks.

VMscape attacks can leak information from a userspace hypervisor to a guest through speculative side-channel techniques.

Values:

- **off**: Disable the mitigation.
- **ibpb** (default): Use the Indirect Branch Prediction Barrier (IBPB) mitigation.
- **force**: Force vulnerability detection even on processors that would otherwise be treated as unaffected.

4.2. UPDATED KERNEL PARAMETERS

mitigations=

[X86, PPC, S390, ARM64, EARLY] Control optional mitigations for CPU vulnerabilities through curated, architecture-independent options.
The **off** option now also disables the **vmscape** mitigation on x86:

- **off**: Disable all optional CPU mitigations. This option improves system performance but can expose users to several CPU vulnerabilities.

The **off** setting is equivalent to the following per-architecture options, if supported:

- **vmscape=off**: Disable VMscape mitigations on x86 systems in addition to previously documented settings.

On x86, after you specify one of the main mitigation modes (such as **off**, **auto**, or **auto,nosmt**), you can additionally use attack-vector based controls as described in **Documentation/admin-guide/hw-vuln/attack_vector_controls.rst**.

rcutree.rcu_normal_wake_from_gp=

[KNL] Reduce the latency of **synchronize_rcu()** calls by maintaining an independent list of callers. This mechanism does not interact with regular RCU callbacks because it does not rely on the **call_rcu()** or **call_rcu_hurry()** paths and applies to normal grace periods only.
You can enable the behavior by writing **1** to **/sys/module/rcutree/parameters/rcu_normal_wake_from_gp** or by passing **rcutree.rcu_normal_wake_from_gp=1** on the kernel command line.

By default, this behavior is now enabled when **num_possible_cpus()** ⇐ **16**, unless you explicitly disable it by passing **rcutree.rcu_normal_wake_from_gp=0** on the kernel command line.

rcutorture.gp_cond=

[KNL] Use conditional or asynchronous update-side normal-grace-period primitives, if available. Previously, this option was documented as using conditional or asynchronous update-side primitives without explicitly clarifying that they apply to normal grace periods.

rh_waived=

Enable waived items in RHEL.

Some specific features or security mitigations can be waived, that is, toggled on or off on demand, in RHEL. However, you should use waivers cautiously, because waiving a mitigation or feature can render a system insecure or even out of scope for support.

Format: **<item-1>,<item-2>,...,<item-n>**

Use the **rh_waived** parameter to enable all waived items that are listed in **Documentation/admin-guide/rh-waived-features.rst**.

4.3. REMOVED KERNEL PARAMETERS

microcode.force_minrev=

[X86] This dedicated parameter for controlling microcode minimal revision enforcement has been removed.
You can now use the **microcode=** parameter with the **force_minrev** option to enable or disable minimal revision enforcement for the runtime microcode loader.

4.4. NEW SYSCTL PARAMETERS

core_sort_vma

The core dump facility now supports sorting virtual memory areas (VMAs) by size in the generated core file.
By default, the kernel writes VMAs in address order. When you set **core_sort_vma** to **1**, the kernel writes VMAs from the smallest size to the largest size. This behavior is known to break at least **elfutils**, but it can be useful when you work with very large or truncated core dumps where the most useful debugging information resides in smaller VMAs.

net.vsock.ns_mode

Control how AF_VSOCK sockets in a network namespace participate in CID allocation and cross-namespace communication.
This setting is read-only. It reports the current namespace's mode, which is determined when the namespace is created and cannot be changed afterwards.

Values:

- **global**: The namespace shares system-wide CID allocation. Sockets in this namespace can communicate with any VM or socket in any namespace that also uses the **global** mode. Sockets in this namespace cannot reach sockets in namespaces that use the **local** mode.
- **local**: The namespace uses private CID allocation. Sockets in this namespace can communicate only with VMs or sockets within the same namespace.

The **init_net** namespace always operates in the **global** mode.

net.vsock.child_ns_mode

Control the default **ns_mode** value for newly created child network namespaces that use AF_VSOCK.
At namespace creation, the kernel initializes **ns_mode** in the child namespace from the parent namespace's **child_ns_mode**. Initially, a namespace's **child_ns_mode** matches its own **ns_mode**.

Values:

- **global**: Child namespaces share system-wide CID allocation. Their VSOCK sockets can communicate with any VM or socket that is reachable from a **global** namespace.
- **local**: Child namespaces use private CID allocation. Their VSOCK sockets can communicate only within their own namespace.

The first write to **child_ns_mode** locks its value. Later writes that set the same value succeed, but writes that attempt to change the value return **-EBUSY**.

Changing **child_ns_mode** affects only namespaces that the kernel creates after the change. Existing namespaces and their children are not modified.

If a namespace runs with **ns_mode=local**, it cannot change **child_ns_mode** to **global**. Attempts to do so fail with **-EPERM**.

4.5. UPDATED SYSCTL PARAMETERS

core_pattern

The **core_pattern** parameter now supports the **%F** specifier to record the pidfd number in core file names.
The following additional format specifier is available:

- **%F**: Insert the pidfd number of the crashing task into the core file name.

CHAPTER 5. DEVICE DRIVERS

Review the new and updated device drivers in Red Hat Enterprise Linux 10.2 to check your hardware compatibility.

5.1. NEW DRIVERS

Table 5.1. Accelerator drivers

Description	Name	Limited to architectures
Driver for Intel NPU (Neural Processing Unit) - 1.0.0	intel_vpu	AMD and Intel 64-bit architectures
Qualcomm Cloud AI Accelerators Accel driver	qaic	AMD and Intel 64-bit architectures

Table 5.2. Bluetooth drivers

Description	Name	Limited to architectures
Bluetooth support for MediaTek devices ver 0.1	btmtk	AMD and Intel 64-bit architectures, 64-bit ARM architecture

Table 5.3. Character device drivers

Description	Name	Limited to architectures
SNP SVSM vTPM (virtual Trusted Platform Module) driver	tpm_svsm	AMD and Intel 64-bit architectures
TPM CRB FFA driver	tpm_crb_ffa	64-bit ARM architecture

Table 5.4. Crypto drivers

Description	Name	Limited to architectures
Intel QuickAssist Technology for GEN6 devices	qat_6xxx	AMD and Intel 64-bit architectures

Table 5.5. Devfreq drivers

Description	Name	Limited to architectures
DEVFREQ userspace governor	governor_userspace	64-bit ARM architecture

Description	Name	Limited to architectures
Generic i.MX bus frequency scaling driver	imx-bus	64-bit ARM architecture
i.MX8M DDR controller frequency driver	imx8m-ddrc	64-bit ARM architecture

Table 5.6. DPLL drivers

Description	Name	Limited to architectures
Microchip ZL3073x core driver	zl3073x	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Microchip ZL3073x I2C driver	zl3073x_i2c	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Microchip ZL3073x SPI driver	zl3073x_spi	AMD and Intel 64-bit architectures, 64-bit ARM architecture

Table 5.7. 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 5.8. Firmware control drivers

Description	Name	Limited to architectures
Firmware control access framework	fwctl	AMD and Intel 64-bit architectures, 64-bit ARM architecture
mlx5 ConnectX firmware control driver	mlx5_fwctl	AMD and Intel 64-bit architectures, 64-bit ARM architecture

Table 5.9. GPIO drivers

Description	Name	Limited to architectures
Intel USBIO GPIO driver	gpio-usbio	AMD and Intel 64-bit architectures

Table 5.10. Graphics drivers and miscellaneous drivers

Description	Name	Limited to architectures
Chrontel ch7006 TV encoder driver	ch7006	AMD and Intel 64-bit architectures
Cirrus driver for QEMU emulated device	cirrus-qemu	AMD and Intel 64-bit architectures
DRM GPUSVM	drm_gpusvm	AMD and Intel 64-bit architectures, 64-bit ARM architecture
DRM GPU scheduler	gpu-sched	IBM Z (s390x)
DRM GPUSVM helper module	drm_gpusvm_helper	AMD and Intel 64-bit architectures
Helpers for DRM sysfb drivers	drm_sysfb_helper	IBM Z (s390x)
Quirks for panel backlight overrides	drm_panel_backlight_quirks	AMD and Intel 64-bit architectures
Silicon Image sil164 TMDS transmitter driver	sil164	AMD and Intel 64-bit architectures

Table 5.11. HID drivers

Description	Name	Limited to architectures
HID driver for Corsair Void headsets	hid-corsair-void	AMD and Intel 64-bit architectures
Intel Touch Host Controller 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 5.12. Hyper-V drivers

Description	Name	Limited to architectures
Microsoft Hyper-V root partition VMM interface	mshv_root	AMD and Intel 64-bit architectures

Table 5.13. I2C drivers

Description	Name	Limited to architectures
Intel USBIO I2C driver	i2c-usbio	AMD and Intel 64-bit architectures
PCA954x I2C multiplexer and switch driver	i2c-mux-pca954x	AMD and Intel 64-bit architectures
Synopsys DesignWare I2C bus adapter in AMD ISP	i2c-designware-amdisp	AMD and Intel 64-bit architectures

Table 5.14. Input drivers

Description	Name	Limited to architectures
Keyboard driver for GPIOs	gpio_keys	AMD and Intel 64-bit architectures, IBM Power Systems (ppc64le)
Polled GPIO buttons driver	gpio_keys_polled	AMD and Intel 64-bit architectures, IBM Power Systems (ppc64le)
Windows-compatible SoC button array driver	soc_button_array	AMD and Intel 64-bit architectures

Table 5.15. Media drivers

Description	Name	Limited to architectures
Conexant cx231xx USB video device driver - 0.0.3	cx231xx	AMD and Intel 64-bit architectures
Conexant CX25840 audio and video decoder driver	cx25840	AMD and Intel 64-bit architectures
cx23415/6/8 driver	cx2341x	AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
Driver for various TV and TV+FM radio tuners	tuner	AMD and Intel 64-bit architectures
Videobuf2 DMA scatter and gather memory handling	videobuf2-dma-sg	AMD and Intel 64-bit architectures
I2C Hauppauge EEPROM decoder driver	tveeprom	AMD and Intel 64-bit architectures
Device node registration for CEC drivers	cec	IBM Z (s390x)
OmniVision OV08X40 sensor driver	ov08x40	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)

Table 5.16. MTD drivers

Description	Name	Limited to architectures
Intel DGFX MTD driver	mtd_intel_dg	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)

Table 5.17. Network drivers

Description	Name	Limited to architectures
Aeonsemi AS21xxx PHY driver	as21xxx	AMD and Intel 64-bit architectures
Intel MLD wireless driver for Linux	iwlmld	AMD and Intel 64-bit architectures, 64-bit ARM architecture
MaxLinear MXL86110 PHY driver	mxl-86110	AMD and Intel 64-bit architectures
Microchip PHY RDS PTP driver	microchip_rds_ptp	AMD and Intel 64-bit architectures
Realtek PHY driver	realtek	AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
Socket CAN driver for Geschwister Schneider and candleLight USB CAN interfaces	gs_usb	AMD and Intel 64-bit architectures, IBM Power Systems (ppc64le)
Common Ethernet library for XDP	libeth_xdp	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Intel Ethernet common library	libie_fwlog	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Intel Ethernet common library admin queue helpers	libie_adminq	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Marvell RVU representor driver	rvu_rep	64-bit ARM architecture
PHY package support	phy_package	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le), IBM Z (s390x)
Realtek 802.11be wireless 8922A driver	rtw89_8922a	AMD and Intel 64-bit architectures, 64-bit ARM architecture
Realtek 802.11be wireless 8922AE and 8922AE-VS driver	rtw89_8922ae	AMD and Intel 64-bit architectures, 64-bit ARM architecture

Table 5.18. Performance monitoring drivers

Description	Name	Limited to architectures
Fujitsu uncore PMU driver	fujitsu_uncore_pmu	64-bit ARM architecture

Table 5.19. Pinctrl drivers

Description	Name	Limited to architectures
AMDISP pin control driver	pinctrl-amdisp	AMD and Intel 64-bit architectures

Table 5.20. Platform drivers

Description	Name	Limited to architectures
AMD 3D V-Cache Performance Optimizer driver	amd_3d_vcache	AMD and Intel 64-bit architectures
AMD HSMP common driver	hsmp_common	AMD and Intel 64-bit architectures
AMD HSMP platform interface driver	amd_hsmp	AMD and Intel 64-bit architectures
AMD HSMP ACPI interface driver	hsmp_acpi	AMD and Intel 64-bit architectures
AMD ISP4 platform driver	amd_isp4	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	intel-oaktrail	AMD and Intel 64-bit architectures
Intel On Demand (SDSi) driver	intel-sdsi	AMD and Intel 64-bit architectures
Intel PMC SSRAM telemetry driver	intel_pmc_ssram_telemetry	AMD and Intel 64-bit architectures
Intel PMT discovery driver	pmt_discovery	AMD and Intel 64-bit architectures
Intel TPMI enumeration module	intel-vsec_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 driver	intel-tpmi_power_domains	AMD and Intel 64-bit architectures

Table 5.21. Thermal drivers

Description	Name	Limited to architectures
Processor thermal PTC interface driver	platform_temperature_control	AMD and Intel 64-bit architectures

Table 5.22. USB drivers

Description	Name	Limited to architectures
Thunderbolt 3 USB Type-C Alternate Mode	typec_thunderbolt	AMD and Intel 64-bit architectures
Intel USBIO bridge driver	usbio	AMD and Intel 64-bit architectures

Table 5.23. vDPA drivers

Description	Name	Limited to architectures
vDPA device in userspace	vduse	AMD and Intel 64-bit architectures, 64-bit ARM architecture

Table 5.24. Virtualization and confidential computing drivers

Description	Name	Limited to architectures
Confidential computing EFI secret area access driver	efi_secret	64-bit ARM architecture

5.2. UPDATED DRIVERS

Table 5.25. Storage driver updates

Description	Name	Current version	Limited to architectures

Description	Name	Current version	Limited to architectures
Broadcom MegaRAID SAS driver	megaraid_sas	07.734.00.00-rc1	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Cisco FCoE HBA driver	fnic	1.8.0.2	AMD and Intel 64-bit architectures
Driver for Microchip Smart Family Controller	smartpqi	2.1.36-026	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
Emulex LightPulse Fibre Channel SCSI driver	lpfc	0:14.4.0.12	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
LSI MPT Fusion SAS 3.0 device driver	mpt3sas	54.100.00.00	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)
MPI3 Storage Controller device driver	mpi3mr	8.15	AMD and Intel 64-bit architectures, 64-bit ARM architecture, IBM Power Systems (ppc64le)

CHAPTER 6. NEW FEATURES AND ENHANCEMENTS

Review new features and enhancements in Red Hat Enterprise Linux 10.2.

6.1. INSTALLER AND IMAGE CREATION

Review new features and enhancements for installer and image creation in Red Hat Enterprise Linux 10.2.

Anaconda supports automatic Flatpak installation from Red Hat Satellite

With this update, Anaconda can automatically install the Flatpak applications during RHEL system installation from Red Hat Satellite. When systems are deployed through Satellite, Anaconda uses the **preinstall.d** mechanism to install Flatpak packages based on the selected environment. For example, the "Server with GUI" environment includes Flatpak-based Mozilla Firefox, ensuring GUI-based systems have necessary applications available immediately after installation. This enables Satellite-managed environments to deliver containerized applications through Flatpak while maintaining existing deployment workflows.

This enhancement ensures Satellite deployments can support RHEL 10 systems with Flatpak-based applications by using familiar installation processes. It also eliminates manual post-installation configuration steps.

Jira:RHEL-95061[1]

Anaconda supports automatic Flatpak installation during system setup

With this update, Anaconda can automatically install Flatpak applications during the RHEL system installation based on the selected environment. This capability works with all installation sources, such as Content Delivery Network (CDN), offline DVD.iso media, and custom LAN servers.

Anaconda installs Flatpak packages by using the **preinstall.d** mechanism during the installation process when users select environments that require Flatpak applications. For example, the "Server with GUI" environment includes Flatpak-based Mozilla Firefox, ensuring GUI-based systems have necessary applications available immediately after installation.

This enhancement enables delivering containerized applications through Flatpak while maintaining a consistent installation experience across all RHEL installation methods. It also eliminates the need for manual Flatpak installation steps after system deployment. You can change the delivery method in Anaconda from Flatpaks back to RPM packages by following the process outlined in the Package selection in Kickstart section of RHEL documentation. For example, use the following configuration to preinstall the Firefox RPM package instead of the Flatpak:

```
%packages
@^graphical-server-environment
-redhat-flatpak-preinstall-firefox
firefox
%end
```

Jira:RHEL-95062[1]

New rdp Kickstart command for remote graphical installation

A new **rdp** Kickstart command was added to enable Remote Desktop Protocol (RDP)-based graphical installations directly from a Kickstart configuration file. The command has the following syntax:

-

```
# rdp [--username <USERNAME>] [--password <PASSWORD>]
```

With this enhancement, you can configure and start a fully automated, headless RDP installation by using Kickstart commands. For complete information about the **rdp** command and its options, see the Kickstart commands reference in the Automatically installing RHEL guide.

Jira:RHEL-96216

Default size for the/boot partition increased to 2 GiB

Before this release, 1 GiB for the /**boot** was often insufficient for systems that require large firmware blobs in **initramfs**. With this update, the default size for the /**boot** partition has been increased from 1 GiB to 2 GiB. This change ensures that there is enough disk space for future kernel updates and associated **initramfs** images. You can manually reduce the partition size or reuse existing smaller partitions when necessary.
Jira:RHEL-151547

RHEL image builder GUI support for creating bootable images

You can create bootable containers and disk images by using the RHEL image builder app in the web console and by using **image-builder-cli**. On first boot, the images automatically subscribe to Red Hat services.

Jira:RHELDOCS-19587[1]

image-builder-cli supports creating PXE boot systems with stateless architecture

With this update, you can use the **image-builder-cli** utility to create stateless PXE images. As a result, you can quickly boot ephemeral nodes that run entirely in RAM over a network by using either an HTTP server or a combined image.
Jira:RHELDOCS-22010

RHEL image builder support for Anaconda network installer images

With this update, you can use RHEL image builder to create Anaconda network installer **.iso** images. By including activation keys directly into the installer, you can automate system registration during the installation process. As a result, instead of standard download pages, you can generate customized, pre-configured images for nightly builds or specific deployment environments.

Jira:RHELDOCS-21852[1]

Finalization locking is available for RHEL on image mode

With this update, you can download bootc system updates without automatically applying them on reboot. You can use the **bootc upgrade --download-only** command to stage updates. To apply the downloaded updates at a later time, use the **bootc upgrade** command. Alternatively, use the **bootc upgrade --from-downloaded** command to apply the staged update without checking the registry for newer versions. The notable enhancements are:

- By staging updates in download-only mode, you can predownload security updates during business hours, validate staged deployments, and choose exactly when to apply them during planned maintenance windows.

- With this feature, you can also ensure better control and security by preventing unintended system updates during routine reboots, and it enables administrators to coordinate controlled rollouts across multiple systems.
- You can apply downloaded updates at any time using bootc upgrade, or you can use **bootc upgrade --from-downloaded** to apply the staged update without checking for newer versions from the registry, which is ideal for scheduled maintenance workflows where the exact downloaded version must be deployed.
- You can apply downloaded updates at any time by using **bootc upgrade**. Alternatively, administrators can use **bootc upgrade --from-downloaded** to apply the staged update without checking the registry for newer versions. This approach is ideal for scheduled maintenance workflows where the exact downloaded version is required for deployment. As a result, operations teams can maintain better governance over production environments, ensuring compliance with strict change control processes, maximizing uptime, and separating network-intensive downloads from actual system changes.

Jira:RHELDOCS-21394[1]

Bootc Virtualization Kit support for bootc

With this update, you can run and convert boot container images into virtual machines. Use the **bcvk** utility to launch ephemeral virtual machines for rapid development and testing, or to generate persistent disk images for production deployments. As a result, your virtual machines run the exact same containerized bootable images used across your environment, maintaining consistency from development to production.

Jira:RHELDOCS-21383[1]

Support for creating stateless PXE images from container builds

You can create stateless PXE images from your container builds in image mode for high-performance computing (HPC) and diskless systems. The build process generates the necessary artifacts, such as **kernel**, **initrd**, and **squashfs**.

Jira:RHELDOCS-20631[1]

6.2. SECURITY

Review new features and enhancements for security in Red Hat Enterprise Linux 10.2.

The system no longer hangs when **fapolicyd** receives SIGSTOP or **ptrace()**

This update of the **fapolicyd-selinux** package introduces an SELinux module to protect the **fapolicyd** service. The new SELinux module prevents users from sending the SIGSTOP signal to **fapolicyd** or tracing **fapolicyd** by using the **ptrace()** function, which might cause the system to crash. As a result, the system no longer hangs or requires manual reboots in the described scenarios.
Jira:RHEL-1368

GSSAPIDelegatedCredentials can be set to **no** in **sshd_config**

With this update, you can set the **GSSAPIDelegatedCredentials** option in the **sshd_config** configuration file to **no**. Although the default value **yes** ensures backward compatibility, you can use **no** for enhanced security control. As a result, an OpenSSH server with **GSSAPIDelegatedCredentials** set to **no** refuses to forward credentials.

Jira:RHEL-5281

New **libreswan-minimal** sub-package reduces container image size

Before this update, the **libreswan** package was a monolithic package with a dependency on **systemd**. This dependency increased the image size of containerized applications.
With this update, the package is modularized by introducing a new **libreswan-minimal** sub-package without dependencies on **systemd** and other optional external tools. As a result, you can create smaller container images for applications that do not use **systemd**. These provide faster startup times and reduced resource usage.

Jira:RHEL-5299

The SELinux policy confines the**redfish-finder** service

New rules in the SELinux policy provide specific confinement for the **redfish-finder** systemd service. This update helps comply with the CIS Server Level 2 benchmark for the restriction of unconfined daemons.
As a result, **redfish-finder** no longer uses the **unconfined_service_t** label and runs correctly in SELinux enforcing mode.

Jira:RHEL-50299[1]

OpenSSH adds support for hybrid ML-KEM NIST

With this update, the OpenSSH suite adds support for the **mlkem768nistp256-sha256** and **mlkem1024nistp384-sha384** key exchange algorithms. As a result, you can protect SSH connections by using the ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism) post-quantum (PQ) key exchange combined with elliptic curves standardized by the National Institute of Standards and Technology (NIST).
Jira:RHEL-70824

libssh supports hybrid key exchange with ML-KEM

With this update, the **libssh** library introduces support for post-quantum traditional (PQ/T) hybrid key exchange methods based on the quantum-resistant Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) standard and traditional Elliptic-curve Diffie-Hellman (ECDH) key exchange schemes. You can use the following methods defined by the Internet Engineering Task Force (IETF) **draft-ietf-sshm-mlkem-hybrid-kex** document in the SSH protocol:

- **mlkem768nistp256-sha256**
- **mlkem768x25519-sha256**
- **mlkem1024nistp384-sha384**

Note that **mlkem768x25519-sha256** is the preferred key exchange method for SSH connections unless you change the configuration.

Jira:RHEL-70825

p11-kit-client.so separates to the**p11-kit-client** subpackage

The **p11-kit-client.so** module moves from the **p11-kit-server** subpackage to the new **p11-kit-client** subpackage. With the separated subpackages, you can install only the required parts and avoid redundant content on host systems or in containers.
Jira:RHEL-89706

OpenSSH relaxed GSSAPI key exchange restrictions in FIPS mode

With this update, the OpenSSH suite permits GSSAPI key exchange methods with the following Diffie-Hellman (DH) and Elliptic Curve Diffie-Hellman (ECDH) groups in FIPS mode:

- **gss-group14-sha256**
- **gss-group16-sha512**
- **gss-nistp256-sha256**

Also, OpenSSH in FIPS mode allows a non-cryptographic use of the MD5 algorithm. As a result, you can use OpenSSH in FIPS mode to establish SSH connections by using GSSAPI key exchange.

Jira:RHEL-91181

New **-c** option for **restorecon** counts relabeled files

With this update, you can use the **restorecon** command with the **-c** option. The **restorecon -c** command performs relabeling, prints the number of relabeled files in its output, and sets the exit code to 0 only if at least one file is relabeled. This makes **restorecon -c** useful for verifying that remediations of labeling problems are successful.
Jira:RHEL-94827

CanonicalMatchUser in **sshd_config** prevents privilege escalation for capitalized AD usernames

This update of the **openssh** packages introduces the **CanonicalMatchUser** directive for the **sshd_config** configuration file. With the new directive, you can configure **Match User** blocks so that **sshd** first attempts to obtain the username from a password database instead of using an alias. As a result, Active Directory (AD) users can no longer bypass chroot restrictions when using capital letters in their usernames, which might lead to privilege escalation.
Jira:RHEL-101440[1]

The SELinux policy confines the **systemd-oomd** service

New rules in the SELinux policy provide specific confinement for the **systemd-oomd** service. This update helps comply with the CIS Server Level 2 benchmark for the restriction of unconfined daemons.
As a result, **systemd-oomd** no longer uses the **unconfined_service_t** label and runs correctly in SELinux enforcing mode.

Jira:RHEL-106998[1]

Several RHEL services transition from SELinux permissive to enforcing mode

With this update, the following SELinux domains move from permissive to enforcing mode:

- **anaconda_generator_t**

- **ktlshd_t**
- **switcheroo_control_t**
- **systemd_pcrextend_t**
- **systemd_user_runtimedir_t**
- **tuned_ppd_t**

These domains temporarily operated in permissive mode. This allowed the system to log additional access denials and gather data to complete the security policy without a service failure. The temporary observation phase is complete.

As a result, the system proactively prevents unauthorized access for these services.

Jira:RHEL-107038[1]

SELinux policy better fits the new OpenSSH structure

With this update, the SELinux policy defines specific security contexts and transitions for the new OpenSSH binary structure, including the **/usr/libexec/openssh/sshd-session** and **/usr/libexec/openssh/sshd-auth** binaries.
The change aligns with splitting the monolithic **sshd** daemon into specialized binaries to reduce the attack surface. By splitting the listener **sshd**, the per-session logic **sshd-session**, and the authentication phase **sshd-auth** into separate processes, the pre-authentication code is isolated in a disjoint address space. This architectural change requires explicit SELinux types to ensure each component maintains the necessary privileges while adhering to the principle of least privilege.

As a result, the OpenSSH server benefits from improved security through process isolation and reduced memory usage after the authentication phase completes. SELinux correctly confines these new binaries, ensuring that host keys and authentication sockets remain protected while allowing standard operations such as PAM authentication to function seamlessly in the new multi-binary environment.

Jira:RHEL-107732

New **setfiles** option reduces memory usage on large file systems

With this update, the **setfiles** utility includes a new **-A** option. Tracking conflicts between inodes with multiple hard links can consume significant memory, especially on large file systems. Use the **-A** option to disable tracking of these conflicts. This reduces memory consumption, allowing to run **setfiles** on memory-constrained systems without encountering high memory overhead.
Jira:RHEL-111505

capnproto is available in the CRB repository

RHEL 10.2 introduces the **capnproto** package, a high-performance data interchange and remote procedure call (RPC) system. This package serves as a shared dependency for **rust-sequoia-sq** and **rust-sequoia-podman**, both of which bundled this library internally before this update.
The **rust-sequoia** packages use the **capnproto** zero-copy serialization and RPC system to communicate with the Sequoia Keystore. This architecture isolates private keys in a separate process to enhance security and ensures the high-speed performance required for large-scale cryptographic tasks, such as container image signing.

The **capnproto** package is available for installation from the CodeReady Builder (CRB) repository. As a result, security updates and bug fixes for the library can be applied independently of the applications that depend on it.

Jira:RHEL-114452[1]

setools rebased to 4.6.0

The **setools** packages, which provide SELinux user-space analysis tools, are rebased to upstream version 4.6.0. This version provides important fixes and enhancements, most notably the following:

- Added the **--role_types** option for the **seinfo** command to display roles allowed for a specified type
- Added a new module to the **sechecker** tool for asserting kernel modules are read-only
- Added support for the **nlmsg** extended permission
- Improved code quality and unit testing
- Dropped methods marked for deprecation

Jira:RHEL-115363

fapolicyd rebased to 1.4.3

The **fapolicyd** packages are rebased to upstream version 1.4.3 and provide many enhancements and bug fixes over the previous version. Most notably:

- Added the **--filter** option for the **fapolicyd-cli --file** command
- Added the **--test-filter** option for the **fapolicy-cli** command to help test filter rules
- Added the **fapolicyd-filter.conf(5)** man page
- Added the **--check-ignore_mounts** option for **fapolicyd-cli**
- Added the **--verbose** flag for the **fapolicyd-cli --check-ignore_mounts** command
- Increased the default value of the **db_max_size** parameter
- Added support for the **db_max_size = auto** option, which enables automatic database size management by the **fapolicyd** daemon
- Increased the default subject cache size
- Moved the **fapolicyd-rpm-loader** program to the **/bin** directory
- Optimized performance of the **fapolicyd** framework

Jira:RHEL-118362

crypto-policies enables ML-KEM for **libssh**

This update of the system-wide cryptographic policies **crypto-policies** adds support for the ML-

KEM (Module-Lattice-Based Key-Encapsulation Mechanism) post-quantum (PQ) key exchange in the **libssh** library. The **mlkem768nistp256-sha256** and **mlkem1024nistp384-sha384** algorithms are enabled by default in all predefined policies. This aligns with support for ML-KEM in OpenSSH, providing a quantum-resistant key exchange method for your SSH sessions.
Jira:RHEL-125889

Support for ML-KEM with NIST curves in FIPS mode added to OpenSSH

This release of the **openssh** packages introduces support for the ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism) post-quantum (PQ) key exchange combined with elliptic curves standardized by the National Institute of Standards and Technology (NIST) in FIPS mode. You can establish SSH connections with a hybrid security provided by the combination of classical cryptography and a quantum-resistant key exchange mechanism.
Jira:RHEL-125929

podman-sequoia supports composite post-quantum signatures

The **podman-sequoia** library provides the ML-DSA-65+Ed25519 and ML-DSA-87+Ed448 algorithms to meet the Commercial National Security Algorithm Suite (CNSA) 2.0 guidelines for software signatures.
As a result, after you install **podman** and **podman-sequoia**, you can create and verify container image signatures with these post-quantum schemes.

Jira:RHEL-126677[1]

capnproto rebased to version 1.3

The **capnproto** package is rebased to version 1.3. This update provides security enhancements and bug fixes, and ensures compatibility with newer Sequoia versions.
Jira:RHEL-127899

/dev/papr-* devices have more specific SELinux labels

With this update of the **selinux-policy** packages, the following devices have more specific SELinux labels:

- **/dev/papr-indices**
- **/dev/papr-physical-attestation**
- **/dev/papr-platform-dump**

This aligns with the addition of new character device interfaces to the kernel, providing user-space application binary interface (ABI) access to the Power Architecture Platform Reference (PAPR) system parameters, in addition to the existing kernel-internal API.

As a result, the SELinux policy assigns distinct labels to these devices so that different permissions can apply to various services accessing them.

Jira:RHEL-129839

libssh rebased to 0.12.0

The **libssh** packages have been upgraded to version 0.12.0. The new version provides many enhancements and bug fixes, notably:

- Added support for hybrid post-quantum key exchange mechanisms, in particular the following:
  - **sntrup761x25519-sha512**
  - **sntrup761x25519-sha512@openssh.com**
  - **mlkem768nistp256-sha256**
  - **mlkem768x25519-sha256**
  - **mlkem1024nistp384-sha384**
- Added support for GSSAPI key exchange as defined in the RFC 4462 and RFC 8732 documents
- Added support for Ed25519 keys through PKCS #11
- Added support for FIDO Universal 2nd Factor (U2F) keys, compatible with OpenSSH
- Added new configuration options:
  - **RequiredRsaSize**
  - **AddressFamily** for clients
  - **GSSAPIKeyExchange**
  - **GSSAPIKexAlgorithms**
- Added more OpenSSH-compatible percent expansion characters
- Added API functions for signing arbitrary data with SSH keys
- Increased the minimum RSA key size to 1024
- Improved the stability and compatibility of the **ProxyJump** directive
- Added functionality to obtain a list of configured identities
- Added new PKI context structure for key operations

Jira:RHEL-133421

crypto-policies enable **mlkem768x25519-sha256** for **libssh**

With this update, the system-wide cryptographic policies enable the **mlkem768x25519-sha256** key exchange algorithm for the **libssh** library in all predefined policies. This aligns with recently added support for this ML-KEM curve hybrid in **libssh**. As a result, **mlkem768x25519-sha256** is enabled by default and negotiated with the highest priority, protecting SSH connections with a combination of traditional and post-quantum cryptography (PQC).
Jira:RHEL-133522

p11-kit rebased to 0.26.1

The **p11-kit** packages have been upgraded to upstream version 0.26.1. The new version provides many enhancements and bug fixes, most notably:

- PKCS #11 headers are updated to version 3.2, which supports post-quantum cryptography (PQC) definitions.
- The trust module now correctly looks up the last DN (Distinguished Name) in the **RDNSequence** attribute as defined in the RFC 4514 document.
- You can specify the server address with the new module configuration option for the Remote Procedure Call (RPC) protocol.
- Handling of an empty array attribute in RPC is fixed.
- Dependency on the **libsystemd** library for server socket activation is removed.

Jira:RHEL-139074[1]

New package: clevis-pin-trustee

The **clevis-pin-trustee** package provides a new Clevis pin **trustee** that enables automated encryption and decryption of LUKS-encrypted volumes by using remote attestation through the Trustee Key Broker Service (KBS). The **trustee** pin integrates with the standard Clevis framework through the **clevis-encrypt-trustee** and **clevis-decrypt-trustee** commands, and it includes a Dracut module **60clevis-pin-trustee** for automated root volume unlocking during early boot.
In scenarios such as confidential clusters for OpenShift and confidential virtual machines with OpenShift Virtualization, the Trustee server acts as the policy enforcement point, releasing the disk encryption key only when the requesting platform's attestation evidence validates against a set of reference values.

As a result, you can bind LUKS-encrypted volumes to one or more Trustee servers by using a **clevis luks bind -d *<device>* trustee '*<config>*'** command. You can also combine the **trustee** pin with other Clevis pins, such as **tang** and **tpm2**, for multi-factor or multi-policy unlock configurations.

Jira:RHEL-139808[1]

Keylime rebased to 7.14.1

The Keylime packages are rebased to upstream version 7.14.1. The most notable bug fixes and enhancements include the following:

Resource management

Resolves a file descriptor leak in the **keylime-policy** tool when processing remote RPM repositories.

Policy tooling

Fixes an issue where the **keylime-policy --ima-measurement-list** option incorrectly handled its default values.

New agent-driven push model

Introduces a new communication mode where agents proactively push attestation data to the verifier. This enhances compatibility with edge computing and environments behind restrictive firewalls or network address translation (NAT).

Expanded hardware cryptography support

Adds support for Elliptic Curve Cryptography (ECC) keys using the P-192, P-224, P-256, P-384, and P-521 NIST curves directly from the Trusted Platform Module (TPM).

Jira:RHEL-140896

keylime-agent rebased to 0.2.9

The **keylime-agent** package is rebased to upstream version 0.2.9, which includes the following enhancements:

New agent-driven push attestation model

The agent supports a push model for attestation. In this model, the agent proactively initiates communication with the verifier rather than waiting for an inbound request. This is particularly beneficial for systems operating behind restrictive firewalls or network address translation (NAT), because it eliminates the need for open inbound ports on the monitored node.

Expanded hardware cryptography support

The agent supports Elliptic Curve Cryptography (ECC) keys generated within the Trusted Platform Module (TPM). Supported NIST curves include P-192, P-224, P-256, P-384, and P-521 to provide more efficient, modern cryptographic operations.

Flexible TPM RSA support

In addition to standard 2048-bit keys, the agent supports alternative RSA key sizes of 1024, 3072, and 4096 bits directly from the TPM. This provides administrators with greater flexibility when aligning with specific organizational security policies or hardware limitations.

ECC-signed TLS certificates

The agent supports using certificates signed with ECC keys for securing TLS communications. This ensures that the entire communication chain between the agent and other Keylime components can utilize high-performance, modern encryption.

Jira:RHEL-140897

crypto-policies enable ML-KEM NIST curves for OpenSSH in FIPS mode

With this update, the system-wide cryptographic policies enable the **mlkem768nistp256-sha256** and **mlkem1024nistp384-sha384** key exchange algorithms for OpenSSH in FIPS mode. This aligns with recently added support for these ML-KEM NIST curve hybrids in OpenSSH. As a result, RHEL 10.2 hosts running in FIPS mode and with the **FIPS** system-wide cryptographic policy active perform SSH key exchanges by using **mlkem768nistp256-sha256** or **mlkem1024nistp384-sha384** as long as the other peer also supports and prefers them.
Jira:RHEL-148560

OpenSCAP rebased to 1.4.3

The OpenSCAP packages have been rebased to upstream version 1.4.3. This version provides bug fixes and various enhancements. For additional information, see the OpenSCAP release notes.
Jira:RHEL-133978

SCAP Security Guide rebased to 0.1.80

For additional information, see the SCAP Security Guide release notes .
Jira:RHEL-152059

6.3. RHEL FOR EDGE

Review new features and enhancements for RHEL for Edge in Red Hat Enterprise Linux 10.2.

The FDO client and servers are fully supported

RHEL 10.2 introduces a new implementation of the FIDO Device Onboarding (FDO) client and servers. These components, which were not available in previous releases, are fully supported and available as the following RPMs:

- go-fdo-client-1.0.0-1.el10.x86_64.rpm
- go-fdo-server-1.0.0-1.el10.x86_64.rpm
- go-fdo-server-manufacturer-1.0.0-1.el10.noarch.rpm
- go-fdo-server-owner-1.0.0-1.el10.noarch.rpm
- go-fdo-server-rendezvous-1.0.0-1.el10.noarch.rpm

WARNING

These new Go-based implementations are not compatible with the original FDO RPMs and container images that remain in Technology Preview. Using the go-fdo-* packages in conjunction with the fdo-* packages or containers is not supported.

Jira:RHELDOCS-18977[1]

The **greenboot-rs** package is available

The **greenboot** health check framework was enhanced as **greenboot-rs**, a reimplementation designed for improved maintainability and supportability. The new version is fully compatible with existing **greenboot** functionality and custom health checks. As a result, this version ensures more robust system roll backs during system upgrades.

Jira:RHELDOCS-21813[1]

6.4. SOFTWARE MANAGEMENT

Review new features and enhancements for software management in Red Hat Enterprise Linux 10.2.

libsolv rebased to 0.7.33

The **libsolv** packages are rebased to upstream version 0.7.33. This version provides the following important fixes and enhancements:

- Removed dependency on the external **find** utility in the **repo2solv** tool.
- Added a new **SOLVER_FLAG_FOCUS_NEW** flag.

- Fixed the return value of the **repodata.add_solv()** function.

Jira:RHEL-86940

librepo rebased to 1.19.0

The **librepo** packages are rebased to upstream version 1.19.0. This version provides the following important fixes and enhancements:

- Fixed a SELinux warning if SELinux runs in a container where **/sys/fs/selinux** is not mounted.
- Fixed caching package checksums on file systems that do not support extended attribute names with uppercase characters.
- When selecting the fastest mirror, mirrors with latency up to twice that of the fastest mirror are randomly shuffled to spread the load.

Jira:RHEL-126292[1]

6.5. SHELLS AND COMMAND-LINE TOOLS

Review new features and enhancements for shells and command-line tools in Red Hat Enterprise Linux 10.2.

Security and TLS improvements in openwsman 2.8.1

The **openwsman** package has been updated to version 2.8.1 with the following improvements:

- Improved TLS 1.3 support.
- Improved compatibility with OpenSSL 3.0.
- Improved SSL/TLS error reporting.
- Improved security by clearing passwords from memory after use and enhancing buffer safety.

Jira:RHEL-99191[1]

opencryptoki rebased to version 3.26.0

The **opencryptoki** packages are rebased to upstream version 3.26.0. This version provides important fixes and enhancements, most notably the following:

- RSA keys up to 16K bits are supported in the **Soft** token and the **p11sak** tool.
- RSA keys up to 8K bits are supported in the **CCA** token. This requires **CCA** version 8.4 or version 7.6 or later.
- The **CKM_SHA512_224_KEY_DERIVATION** and **CKM_SHA512_256_KEY_DERIVATION** key derivation mechanisms are supported in the **Soft** and **ICA** tokens.
- The **CKK_SHAxxx_HMAC** key types and **CKM_SHAxxx_KEY_GEN** key generation mechanisms are supported in the **Soft**, **ICA**, **CCA**, and **EP11** tokens, as well as the **p11sak** tool.

- Key wrap and unwrap commands to export and import private and secret keys by using various key wrapping mechanisms are supported in the **p11sak** tool.
- Using a hardware security module (HSM)-protected TLS client key through a PKCS #11 provider is supported in **p11kmip**.
- Exporting non-sensitive private keys to password-protected PEM files is supported in the **p11sak** tool.
- Canceling an operation by using a **NULL** mechanism pointer at the **C_XxxInit()** call is supported as an alternative to **C_SessionCancel()** for PKCS#11 version 3.0.
- Pairing the friendly BLS12-381 elliptic curve (EC) for sign and verify operations by using **CKM_IBM_ECDSA_OTHER** and signature and public key aggregation by using **CKM_IBM_EC_AGGREGATE** is supported in the **EP11** token.
- Generating BLS12-381 EC keys is supported in **p11sak**.
- IBM-specific ML-DSA and ML-KEM key types and mechanisms are supported in the **EP11**, **CCA**, and **Soft** tokens, and **p11sak**. Before you use these key types and mechanisms, note the following requirements:
  - The **EP11** token requires an **EP11** host library version 4.2 or later, and a CEX8P cryptographic card with firmware version 9.6 or later on IBM z17 or version 8.39 or later on IBM z16.
  - The **CCA** token requires **CCA** version 8.4 or later.
  - The **Soft** token requires OpenSSL 3.5 or later, or a configured OQS-provider.

Jira:RHEL-100058[1]

Overriding the systemd-logind session class for cron-initiated sessions

With this update, you can override the **systemd-logind** session class for sessions that **cron** scripts start. To start a session without triggering the **systemd --user** manager, set the **XDG_SESSION_CLASS=background-light** environment variable in the crontab. This configuration reduces the number of log messages that **cron** executions generate.

Jira:RHEL-109832[1]

Environment modules rebased to v5.6.1

Environment modules is rebased to upstream version 5.6.1. This release introduces key new features, enhancements, documentation and community updates, and few bug fixes. Here is the list for reference:

- New features and enhancements:
  - Recursive module searching with spider command: With this update, you can use the **spider** sub-command to find available modules in enabled modulepaths and recursively within modulepaths enabled by those modules. You can control the output depth and content by using the **--indepth** switch or the **spider_output** configuration option.

    - Module aliases with provide command: This enhancement introduces the **provide** modulefile command, which defines an alias for the currently evaluated module and communicates when a module offers additional components or functionality.
    - Automated conflict handling: With this update, the **conflict_unload** configuration option automatically unloads conflicting modules and their dependents when you load a new module. You must enable both **auto_handling** and **conflict_unload** to activate this automated behavior.
    - Integrated information logging: This update adds logging capabilities through the **logger** and **logged_events** configuration options. You can now track module commands and evaluations in the system log.
    - Module help and warning commands: This release introduces the **module-help** command to define help text for modules and the **module-warn** command to issue warnings when a module is evaluated.
- Bug fixes
    - Path resolution in modulefile commands: Before this update, the behavior of path resolution was unclear. This release clarifies that no automatic path resolution is performed on **prepend-path**, **append-path**, or **remove-path** commands. For detailed information about changes, refer to the Environment Modules upstream documentation .

Jira:RHEL-132336

6.6. INFRASTRUCTURE SERVICES

Review new features and enhancements for infrastructure services in Red Hat Enterprise Linux 10.2.

foomatic-rip filter rejects unrecognized PPD values

The **foomatic-rip** filter rejects PostScript Printer Description (PPD) values not in an approved list of hashes. Before this update, certain PPD options were vulnerable to security exploits. This update implements an allowlist mechanism to ensure secure printing.
For new installations, use the **foomatic-hash** tool to scan the PPD file and move approved hashes to the **/etc/foomatic/hashes.d/** directory. For existing installations, review auto-allowed values in the **/var/tmp/foomatic.*** file.

Jira:RHEL-93944[1]

PHP 8.4 available

RHEL 10.2 provides PHP in version 8.4. This version provides many enhancements and bug fixes over version 8.3, most notably:

- Property hooks provide support for computed properties natively understood by IDEs and static analysis tools.
- Asymmetric visibility controls the scope to write to a property independently from the scope to read the property.
- The **#[\Deprecated]** attribute makes the existing deprecation mechanism available to user-defined functions, methods, and class constants.

- A new DOM API is available within the **Dom** namespace, which includes standards-compliant support for parsing HTML5 documents.
- The **BcMath\Number** object enables object-oriented usage and standard mathematical operators when working with arbitrary precision numbers.
- The **array_find()**, **array_find_key()**, **array_any()**, and **array_all()** functions are available.
- You can access properties and methods of a newly instantiated object without wrapping the **new** expression in parentheses.

Jira:RHEL-105827[1]

chrony rebased to version 4.8

The **chrony** packages are rebased to upstream version 4.8, which includes the following notable enhancements and bug fixes:

- The **maxunreach** option is added to limit the selection of unreachable sources.
- The **-u** option is added to the **chronyc** command to drop root privileges.
- The **opencommands** directive is added to select remote monitoring commands.
- The **waitsynced** and **waitunsynced** options are added to the **local** directive.
- The RTC **refclock** driver is added.
- You can specify the PHC **refclock** driver with a network interface name.
- Detection of clock interference from other processes is added.
- The **chronyc** socket is hidden to mitigate unsafe permissions changes.
- The **refclock** samples are validated for reachability updates.

Jira:RHEL-112593

valgrind rebased to upstream version 3.26.0

The upgrade to the upstream version 3.26.0 provides the following notable enhancements:

- valgrind recognizes the following Linux kernel system calls: **cachestat**, **futex_waitv**, **listmount**, **mount_setattr**, **mseal**, **quotactl_fd**, **remap_file_pages**, **setdomainname**, **statmount**, **swapoff**, **swapon**, **sysfs**, and **ustat**.
- A new option, **--modify-fds=yes**, has been added. This option behaves like **--modify-fds=high**, returning the highest available file descriptor first. However, if file descriptors **0**, **1**, or **2** (**stdin**, **stdout**, **stderr**) are available, they are returned before higher-numbered file descriptors.
- When **--xml=yes** is used, log output protocol version 6 is always enabled. Protocol version 6 includes error summaries in the XML output.

- A new value, **bad**, has been added for the **--track-fds** option. When **--track-fds=bad** is specified, valgrind reports only invalid file descriptor usage, such as double close or use of an invalid file descriptor. It does not report unclosed file descriptors at program exit.
- DWARF inlined subroutine handling has been rewritten to work across compilation units. This update removes backtraces that previously displayed **UnknownInlinedFun** in warnings or error messages.
- A new utility script, **vgstack**, has been added. Use **vgstack <PID>** to attach to a running valgrind process and display backtraces of the target executable. The script provides the following options:
  - **-h** - Displays minimal help.
  - **-v** - Displays version information.

Jira:RHEL-120966

SystemTap is rebased to version 5.4

SystemTap is rebased to version 5.4. The notable changes in this update include:

- **Implicit Header Discovery**: The **@cast()** operator now automatically searches the Linux Userspace API (UAPI) **<vmlinux.h>** header for type declarations. This reduces the requirement for manual header file inclusion in many common tracing scenarios.
- **Enhanced Type Validation**: Improvements to type checking and autocast processing provide more rigorous analysis during the translation phase, identifying potential type mismatches earlier in the development cycle.

Jira:RHEL-121663

elfutils rebased to 0.194

The upgrade to the upstream version 0.194 provides the following notable enhancements:

- **debuginfod-find**: Fixed a caching issue that prevented re-downloading files after a user-cancelled download.
- **elfclassify**: Added the following new options:
  - **--has-debug-sections**
  - **--any-ar-member**
- **elflint**: Vendor and application-specific ELF note types no longer trigger compliance errors.
- **libdwfl_stacktrace**: Added a new function, **dwflst_sample_getframes**.
- **libelf**: Added manual pages for many library functions.
- **readelf**: Improved performance by up to 13% when using the **-N** option.

Jira:RHEL-121665

sscg rebased to version 4.0.3

The **sscg** packages are rebased to upstream version 4.0.3. This version provides important fixes and enhancements, most notably the following:

- Module-Lattice-Based Digital Signature Algorithm (ML-DSA) key generation is supported to provide post-quantum cryptography capabilities.
- Elliptic Curve Digital Signature Algorithm (ECDSA) key generation is supported.
- The command-line interface help output is reorganized into logical groups.

Jira:RHEL-123675

Apache's **ErrorLogFormat** supports millisecond timestamps

With this update, Apache's **ErrorLogFormat** supports millisecond timestamps. Millisecond-level timestamps in error logs improve log filtering, troubleshooting efficiency, and cross-system traceability. You can configure this, for example, by using the **%{m}t** format specifier. As a result, you can correlate and filter logs across systems with millisecond precision.

Jira:RHEL-145713[1]

6.7. NETWORKING

Review new features and enhancements for networking in Red Hat Enterprise Linux 10.2.

Nmstate can configure Libreswan and use its default values

By default, the NMstate API uses NetworkManager to send configurations to Libreswan service. In this case, NetworkManager defines default values, which are different from Libreswan's defaults. With this enhancement, you can set **nm-auto-defaults: false** in the YAML file and Nmstate does not inject any extra settings. In this case, Libreswan uses this configuration and also its own default values.

For backward compatibility, the default value of **nm-auto-defaults** is **true**.

Jira:RHEL-26350

The NetworkManager Libreswan plugin and Nmstate support using a single tunnel for multiple subnets

This update enhances the NetworkManager Libreswan client plugin and Nmstate to configure multiple subnets in IPsec policies. This corresponds to the use of multiple subnets in the **leftsubnets** and **rightsubnets** parameters in the Libreswan configuration. As a result, users can connect to multiple subnets by using a single IPsec tunnel.

Jira:RHEL-33712

NetworkManager-libreswan supports on-demand IPsec connections

With this enhancement, you can use the **NetworkManager-libreswan** plugin to start Libreswan IPsec connections in listening mode. Previously, NetworkManager failed to activate a connection if the remote endpoint was unreachable. By setting the new **nm-connect-mode** property to **ondemand** in the connection profile, the tunnel remains active in a listening state after an initial failure. This ensures the system can still accept incoming connection requests even if it could not initiate the primary tunnel.

Jira:RHEL-67307

The **epoll** kernel API supports IRQ suspension for improved network efficiency

This enhancement adds IRQ suspension support to the **epoll** kernel API. This improves network processing efficiency within the kernel stack. This mechanism bridges the gap between throughput and latency by providing a way to dynamically optimize the networking stack for high-load efficiency and low-load responsiveness simultaneously. Applications that use **epoll** with this new mechanism can reduce CPU cycle consumption during high traffic loads and decrease tail latency during low traffic periods.
Note that you must modify your application to support this IRQ suspending.

Jira:RHEL-77189[1]

Nmstate can set alternative names on network interfaces

With this enhancement, you can use the Nmstate API to set alternative names on network interfaces to simplify configuration management and support processes. For example, to assign **LAN** as an alternative name to **enp1s0** and remove the name **internal-LAN**, use:

```
interfaces:
  - name: enp1s0
    alt-names:
      - name: LAN
      - name: internal-LAN
        state: absent
```

Jira:RHEL-90096

iproute rebased to version 6.17.0

The **iproute** package has been updated to upstream version 6.17.0.
Notable enhancements:

- The **tc** utility supports 64-bit hardware packet counters.
- The **ip** utility displays the **netns-immutable** property.
- The **ip** utility supports the **IFLA_VXLAN_MC_ROUTE** configuration attribute.
- The **ip neigh** command supports the **extern_valid** flag.
- The **ip rule** command supports port and Differentiated Services Code Point (DSCP) mask.
- The **ip stats** command supports bridge VLAN statistics.
- The **bridge fdb** command supports the forward database (FDB) activity notification control.
- The **bridge mdb** command supports the offload failed flag.
- The color output handling was improved.

Jira:RHEL-98263

NetworkManager supports specifying an HSR interlink interface

With this update, RHEL users can configure an interlink interface for High-availability Seamless Redundancy (HSR) connections. Users can now use the **hsr.interlink** property to specify the interlink interface name. As a result, you can configure RHEL as a Redundancy Box (RedBox).
Jira:RHEL-100768

The PRP and HSR protocols are fully supported

The **hsr** kernel module 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.

  The protocols were previously available as a Technology Preview. Starting with RHEL 10.2, Red Hat fully supports this module.

Jira:RHEL-100942[1]

Setting the DHCP client ID is now possible through a kernel argument

With this update, users can now set the DHCP client ID as a kernel argument. Certain DHCP servers require this ID to identify a client correctly. By setting the **rd.net.dhcp.client-id** kernel argument, the client ID is already available during early boot operations.
Jira:RHEL-108454

The **nftables** framework supports name-based **netdev** hooks with wildcards

This enhancement introduces name-based **netdev** hooks with wildcard support to the **nftables** kernel component. This ensures defined rule sets remain stable regardless of interface presence. Previously, **nftables** would bind to each specified interface immediately upon adding a **flowtable** or **netdev**-family chain. Consequently, the transaction failed due to a non-existing interface, and removing an interface deleted the matching interface specifications or entire bound chains.
With this update, hooks for non-existing interfaces are accepted in an inactive state and bind to matching interfaces at the time they appear in the system. This dynamic registration also provides the possibility to accept simple interface (suffix) wildcards to bind a **flowtable** or **netdev**-family chain to any matching interface. You can inspect currently active hooks by using the **nft list hooks** command.

Jira:RHEL-108861

RHEL supports WiFi7 hardware

RHEL 9.8 added support for WiFi7 hardware. You can use use it to connect your host to wireless networks that use this standard.

Jira:RHEL-111098[1]

The kernel supports setting a lower TCP maximum retransmission timeout value

With this enhancement, you can set a lower maximum TCP retransmission timeout value than the default **120000** ms to reduce network latency. Note that changing this setting can require tuning other kernel settings as well.

You can configure this limit either through the **tcp_rto_max_ms** kernel **sysctl** setting or the **TCP_RTO_MAX_MS** socket option. If you set both, the socket option has a higher priority.

Jira:RHEL-115393[1]

FRR rebased to version 10.4.1

The FRR is now rebased to version 10.4.1. This version fixes several issues affecting stability, correctness, and reliability. Notable changes include:

- BGP (**bgpd**):
  - Resolved **addpath** handling issues that could incorrectly withdraw selected routes.
  - Fixed link-local next-hop capability handling.
  - Corrected a compilation issue in the **bgpd** module.
  - Improved graceful restart behavior by fixing the **selectionDeferralTimer** display.
  - Addressed initialization issues with local variables.
  - Reversed changes related to EVPN testing that caused instability in non-default EVPN backbone configurations.
- OSPF (**ospfd**):
  - Fixed a use-after-free issue related to LSA handling, improving daemon stability. EIGRP (**eigrpd**):
  - Improved validation of hello packets and TLVs to enhance protocol robustness.
- Zebra and core libraries:
  - Fixed buffer overflow issues identified through fuzz testing.
  - Improved handling of singleton nexthops during link state changes for weighted ECMP (WCMP).
  - Corrected computation of link-state ZAPI message sizes.
- VTYSH:
  - Fixed an issue where copying configuration from a file did not correctly apply settings.
- Testing improvements:
  - Enhanced reliability of embedded route processor (RP) topotests.
    These updates improve overall routing stability, correctness, and resilience, particularly in dynamic or large-scale network environments.

Jira:RHEL-118620

nftables rebased to version 1.1.5

The **nftables** package has been updated to upstream version 1.1.5.
Notable enhancements:

- The memory consumption with sets and maps was reduced.
- You can use protocol dependency values in sets.
- The auto-merge feature skips elements with timeout and expiration.
- You can use the **queue** keyword in set type definitions by using the **typeof** keyword.
- The **nft monitor** command can monitor **flowtable** events.
- For consistency with other commands, the **nft list sets inet *<table_name>*** command works without the **table** keyword.
- The **nftables** framework internally uses a range expression to represent a range instead of two comparisons.
- A symbol table for Multipath TCP subtypes was added. With this feature, you no longer need to look up actual subtype values in the respective RFC.
- Support for mangling **bitfield** headers was added.
- Set elements with multi-word descriptions are now displayed in a single line.
- The layer 4 protocol dependency when listing raw expressions is no longer removed.
- The JSON format supports the **typeof** keyword.
- The bytecode generation for Virtual Local Area Network (VLAN) Priority Code Point (PCP) mangling in **netdev**-family chains was fixed.
- An issue causing bogus elements in large concatenated set ranges was fixed.
- A new check result was added to the Forwarding Information Base (FIB) expression to verify routes.
- The total number of elements is now displayed when listing sets.
- You can delete maps by using their unique handle.
- The JSON parser was hardened.

Notable bug fixes:

- Error messages for set or map re-declarations with conflicting types were improved.
- The **optimize** parameter was fixed and improved.
- Extended error reporting with large set elements was fixed.
- **nftables** avoids the incorrect removal of **meta nfproto** matches in listings.
- The **get** and **reset** commands with interval sets and maps were fixed.
- Device names in **basechain** and **flowtable** declarations are quoted.
- A misleading **No buffer space available** error message was corrected.

Jira:RHEL-121194

VLAN segmentation support for HSR and PRP interfaces

With this enhancement, you can create VLAN interfaces on top of High-availability Seamless Redundancy (HSR) and Parallel Redundancy Protocol (PRP) interfaces to enable network traffic segmentation. When configured, the kernel adds a VLAN tag to all packets transmitted through the VLAN interface. This provides greater control over traffic isolation. Note that supervision frames remain unaffected by this configuration and are always transmitted without a VLAN tag.

Jira:RHEL-130475[1]

Threaded NAPI busy polling support added

With this enhancement, you can enable threaded NAPI busy polling on RHEL to significantly reduce the network latency. This feature uses dedicated kernel threads to continuously check for incoming packets, rather than waiting for hardware interrupts.
By design, threaded NAPI busy polling consumes more CPU cycles to achieve higher performance and lower latency, as the CPU remains active to process data immediately. Threaded NAPI busy polling is beneficial for high-performance, low latency applications, for example applications that use the **AF_XDP** socket. Use this enhancement for workloads where predictable, sub-microsecond response times are critical.

Jira:RHEL-130765[1]

The dpll utility can manage and monitor DPLL devices

With this update, the **iproute** package includes the **dpll** utility which you can use to manage and monitor digital phase-locked loop (DPLL) devices. The utility uses **libmnl** to communicate with the kernel through the **netlink** interface, providing a configuration tool for DPLL devices and pins.
Jira:RHEL-131660

The K1 power state flag can be disabled one1000e NICs

The K1 state reduces power consumption on ICH-family network interface controllers (NIC) during idle periods. However, on Intel Meteor Lake and later platforms, enabling K1 state on NICs that use the **e1000e** driver can cause packet loss due to firmware misconfiguration, interoperability with certain link partners, and other conditions.
Default:

- The K1 state is disabled on Intel Meteor Lake and later platforms.
- The K1 state is enabled on platforms earlier than Intel Meteor Lake.
  If you experience problems related to the K1 power state, disable K1 for the affected device:

  1. Display the current status:

     ```
     # ethtool --show-priv-flags <device>
     ...
     disable-k1: off
     ```

  2. Disable the K1 state:

     ```
     # ethtool --set-priv-flags <device> disable-k1 on
     ```

Jira:RHEL-134991[1]

Qualcomm wireless cards work correctly if passed through to a VM

Due to missing upstream support for passing Qualcomm wireless cards to VMs by using the PCI pass through feature, these cards do not work correctly in VMs. With this update, the **ath11k** and **ath12k** drivers use certain kernel parameters to work around the problem. As a result, Qualcomm wireless cards that use these drivers work if you pass the devices to VMs. Note that the solution is only an unsupported workaround.

Jira:RHEL-141347[1]

The FOU and GUE protocols added to the kernel

This update adds the **fou** and **fou6** modules to the **kernel-modules-extra** package. With these modules, you can configure connections that use the following protocols:

- Foo-over-UDP (FOU), which encapsulates IP protocols directly within UDP packages, without adding extra headers. For example, you can use this protocol for tunneling protocols, such as Generic Routing Encapsulation (GRE) or IP-in-IP (IPIP).
- Generic UDP Encapsulation (GUE), which adds a small header inside the UDP payload to carry metadata, such as the inner protocol. With GUE, you can use multiple protocols on the same UDP port, which makes GUE more flexible than FOU.
  Red Hat does not support the **fou** and **fou6** kernel modules.

Jira:RHEL-142435[1]

The **firewalld** service supports policy sets

This enhancements adds support for policy sets to the **firewalld** service. Policy sets are collections of pre-defined policies that you can use as a starting point for certain configurations. For example, the **gateway** policy set is a set of configurations that enable masquerading, connection tracking helpers, and forwarding between zones.
For further details, see Using **firewalld** policy sets to configure a router .

Jira:RHEL-70357[1]

6.8. KERNEL

Red Hat Enterprise Linux 10.2 is distributed with the kernel version 6.12.0-211.7.1. Review new features and enhancements for kernel in Red Hat Enterprise Linux 10.2.

The **perf** command supports Monaka CPU for performance monitoring

With this update, the **perf** command supports Monaka CPU to enable performance monitoring in the system. As a result, you can use this feature to analyze performance and power for high-performance computing (HPC) and datacenter applications. This feature is integrated into the Linux kernel version 6.12.0 and later.

Jira:RHEL-23107[1]

LUKS volume key for secure**vmcore** data saving on RHEL systems

With this update, you can pass the LUKS volume key to the **kdump** kernel, to save **vmcore** data to a

LUKS-encrypted disk volume. This enhancement secures **vmcore** data on RHEL systems, as sensitive data remains protected in the event of system crashes. To activate this optional feature, you must use the 'kdumpctl setup-crypttab' command. This update is available for the x86_64 architecture in RHEL 10.2.
Jira:RHEL-29037

PerfMon support added for Clearwater Forest on CentOS Stream kernel

With this update, PerfMon support is added for Clearwater Forest, a hardware or software platform, on the CentOS Stream kernel. This enhancement enables performance monitoring for the Clearwater Forest platform, improving overall system efficiency and stability.

Jira:RHEL-45066[1]

EDAC Driver Adds Intel Clearwater Forest Server Support

The EDAC driver is updated to add platform support for Intel Clearwater Forest (CWF) servers, enhancing RAS capabilities for this hardware. This change improves error detection and correction functionality specific to the Intel platform.

Jira:RHEL-45084[1]

Perf tool rebased to upstream version 6.17

The perf tool and its kernel backend are rebased to align with upstream version 6.17. This update introduces several enhancements and bug fixes. Most notably, the following:

- Addressed memory leaks in perf trace.
- Supports the RDPMC metrics in clear mode.
- Added RAPL energy events support in the perf tool for the ARL-U platform. These changes improve performance analysis and resolve known issues in the perf tool.

Jira:RHEL-78200[1]

bpf is rebased to version 6.17

- The eBPF subsystem is rebased to the Linux kernel upstream version v6.17. This update includes the following changes and enhancements:
- New eBPF kernel functions (**kfuncs**):
  - **bpf_cpumask_populate()** for populating CPU mask bits
  - **bpf_copy_from_user_task_str()** for reading strings from another process address space
  - **bpf_dynptr_copy()** for copying dynamic eBPF pointers
  - **bpf_set_dentry_xattr()** and **bpf_remove_dentry_xattr()** to set and remove xattrs with the **security.bpf.** prefix
  - **bpf_rbtree_left()**, **bpf_rbtree_right()**, and **bpf_rbtree_root()** for traversing the eBPF rbtree data structure

    - functions for reading memory into eBPF dynamic pointers
    - functions for read-only string operations
    - **bpf_cgroup_read_xattr()** to read **xattr** of a cgroup node
- Improved verification of eBPF programs with loops
- Referenced pointers (**kptrs**) can now be passed into **struct_ops** callbacks
- Reduced **bpf_cgrp_storage_busy()** false positives when accessing cgroup local storage
- New mechanisms for ordering of cgroup eBPF programs
- The eBPF Token can now delegate the privilege to read BTF data to user-space applications
- It is no longer possible to attach eBPF programs to **noreturn** functions
- New locking mechanism, Resilient Queued Spinlock, which makes BPF programs and map operations less likely to deadlock the running kernel
- Support for up to 12 arguments in BPF trampoline on **arm64**, enabling eBPF trampolines for kernel functions with more than 7 parameters
- Support for **mmap** of **vmlinux** BTF data
- New eBPF iterator for traversing the list of all DMA buffers
- eBPF streams for error reporting of various conditions detected by the eBPF runtime
- Improved precision for **BPF_ADD** and **BPF_SUB** operations in the verifier
- Support for calls to **bpf_rdonly_cast(v, 0)** that logically correspond to casts to **void ***
- Support for new eBPF instructions: **load_acquire**, **store_release**, and timed **may_goto**
- Support for atomic update of eBPF maps that contain a hash table of eBPF maps
- Method for retrieving file descriptor information for eBPF links

Jira:RHEL-78204[1]

Perf tool rebased to upstream v6.18

The perf tool and its kernel backend are rebased to align with upstream version v6.18. This update introduces several enhancements and bug fixes. Most notably, the following:

- Addressed memory leaks in perf trace.
- Supports the RDPMC metrics in clear mode.
- Added RAPL energy events support in the perf tool for the ARL-U platform. These changes improve performance analysis and resolve known issues in the perf tool

Jira:RHEL-78292[1]

cpupower Python bindings are now in RHEL 10**kernel-tools-libs-devel**

With this update, the **cpupower** Python bindings are integrated in RHEL 10. This enhancement places the bindings in the **kernel-tools-libs-devel** package for easier access.

Jira:RHEL-83442[1]

Userspace action triggers for rtla

With this update, the **rtla** tool now supports triggering userspace actions either when a latency threshold is reached or tracing concludes. With **rtla**, you can execute diagnostic commands or extract trace data before the instance is removed, regardless of whether a threshold violation occurred.

Jira:RHEL-89807[1]

Intel QAT GEN6 device driver support

The Intel QAT crypto device driver is updated to support QAT GEN6 devices through the new **qat_6xxx** driver. GEN6 devices enable concurrent use of symmetric encryption, asymmetric encryption, and data compression. This was not available in earlier generations.

Jira:RHEL-94928[1]

tpm2-tools rebased for TPM 2.0 improvements

The **tpm2-tools** package is updated to ensure compatibility with modern TPM 2.0 hardware and improve security tooling support. This update enables enhanced TPM-based operations and aligns with upstream security and feature developments.

Jira:RHEL-94930[1]

Device IDs are added for the In-memory Analytics Accelerator (IAA) on the Wildcat Lake platform

With this update, the IAA is now moved from a Technology Preview to the supported state and the device IDs are added for In-memory Analytics Accelerator (IAA). As a result, devices on the Wildcat Lake platform are now supported.

Jira:RHEL-95628[1]

Enhanced kernel issue debugging with the**function_graph** tracer on RHEL

With this update, you can trace and debug kernel issues more effectively on Red Hat Enterprise Linux (RHEL). This feature displays return values of functions within the function graph by using the **function_graph** tracer in **ftrace**. As a result, debugging experience improves for developers and system administrators.

Jira:RHEL-105766[1]

View CVEs patched by live kernel updates

kpatch reports which kernel CVEs are patched by live patches for the currently running base kernel. This enhancement helps administrators verify that specific CVEs are already remediated through live patching even when the on-disk kernel version appears vulnerable.
By listing CVEs that are patched only by **kpatch**, this enhancement improves security reporting and enables integration with compliance workflows and external scanners that must account for live-patched vulnerabilities.

Jira:RHEL-106283[1]

Updating **kernel** CCP crypto driver support for Venice PCI device

This update adds support for the AMD Venice CCP crypto device with PCI device ID 0x17D8 (PCIID 1002:17D8) in the kernel CCP driver. This enables systems with the Venice CCP hardware to use the updated cryptographic offload capabilities provided by the device.

Jira:RHEL-106909[1]

crash rebased to 9.0.1

The **crash** package, which provides a kernel analysis utility for live systems and various types of dump files, is rebased to upstream version 9.0.1. This version provides a number of fixes and enhancements, most notably the following:

- Internal **gdb** is updated to version 16.2.
- Added **gdb multi-stack** unwind support on 64-bit architectures (x86-64-v3), aarch64, and ppc64.
- Added Rust support.

Jira:RHEL-114659

You can select **cyclictest** or **timerlat** as the measurement modules in **rteval**

With this update, you can select the measurement module for the **rteval** utility. This overrides the default setting in the **rteval.conf** file. This new feature, 'measurement-module', provides greater flexibility and control over performance testing, which enhances the precision and customization.

Jira:RHEL-114927[1]

Optimize CPU usage with Tuna 10.2's **libcpupower** functionality

With this update, you can manage CPU idle states more effectively in Tuna 10.2. The **libcpupower** functionality has been re-enabled, which allows disabling, enabling, or checking the status of idle states on selected CPUs. By using the **tuna cpu_power** command, you can optimize your CPU usage.

Jira:RHEL-116084

6.9. BOOT LOADER

Review new features and enhancements for boot loader in Red Hat Enterprise Linux 10.2.

Support for Dynamic Key Management in PowerVM LPAR Secure Boot (GRUB2) on IBM Power Systems

With this release, PowerVM LPAR guest operating systems on IBM Power Systems support dynamic key management for secure boot verification. This enhancement allows you to enroll and manage your own keys in the Platform Key Store, transitioning from a static key model.

During boot, the partition firmware authenticates **grub2** using the enrolled verification key. Then **grub2** verifies the kernel image integrity before loading. This improves flexibility and control over boot integrity and strengthens the security posture for IBM Power Systems environments.

Jira:RHEL-24510[1]

BLS snippets support the efi keyword for UKI

You can create Boot Loader Specification (BLS) snippets for kernel unified kernel images (UKIs) and use the **efi** keyword to specify the path to the UKI, similar to how the **linux** keyword specifies the path to the kernel. For example:

```
title Red Hat Enterprise Linux 10.2 (6.12.0-197.el10)
version 6.12.0-197.el10.x86_64
efi /EFI/Linux/kernel-6.12.0-197.el10-UKI.efi
```

In this configuration, BLS snippets reside in **/boot/efi/loader/entries**, and the UKIs reside in **/boot/efi/EFI/Linux**.

Jira:RHEL-119685

shim signed with Microsoft 2011 and 2023 UEFI certificates

The **shim** bootloader package is signed with both the Microsoft Windows UEFI Driver Publisher (MS 2011) certificate and the Microsoft UEFI CA 2023 certificate for Red Hat Enterprise Linux 10.2. This update helps maintain compatibility with systems that rely on either of these Microsoft UEFI trust anchors while preserving the existing Red Hat UEFI Publisher 2024 signature.
With this change, both **shimx64.efi** and **shimaa64.efi** binaries are correctly signed, enabling secure boot environments to validate the updated bootloader components on supported hardware platforms.

Jira:RHEL-144033

6.10. FILE SYSTEMS AND STORAGE

Review new features and enhancements for file systems and storage in Red Hat Enterprise Linux 10.2.

A new watchdog for fanotify permission events is now available

With this update, an optional watchdog for **fanotify** permission events has been introduced. If a system hang occurs due to **fanotify** permission events, the watchdog logs the process ID and name of the task responsible for the hang to the system log. This enhancement simplifies and accelerates the diagnosis of **fanotify** related hangs without requiring kernel crash dump analysis.
Note that the watchdog is disabled by default. To enable it, write a timeout value to **/proc/sys/fs/fanotify/watchdog_timeout**. When enabled, the watchdog incurs negligible performance overhead.

Jira:RHEL-44601[1]

LVM now supports Persistent Reservations on volume groups

With this update, the Logical Volume Manager (LVM) has been enhanced to manage persistent reservations on a volume group (VG). With this feature, LVM controls access and ownership of shared storage resources used by Volume Groups. This can be useful in clustered environments that use shared block storage. For more information, see the **lvmpersist(8)** man page on your system.
Jira:RHEL-60931

io_uring interface added for asynchronous I/O

The **io_uring** interface supports asynchronous I/O operations. With this update, applications use this

interface to submit multiple I/O requests without blocking the calling process. **io_uring** uses shared ring buffers between user space and kernel space to reduce system call overhead and avoid buffer copying. This interface is more efficient and supports more asynchronous system calls than Linux AIO.

Jira:RHEL-120700[1]

Stratis now maintains volume keys in the process keyring for encrypted pools

With the release of **stratisd 3.8.6** and **stratis-cli 3.8.3**, the Stratis storage management system can now automatically maintain the volume keys of encrypted pools.
Previously, if **stratisd** needed to extend an encrypted pool automatically, the operation could fail if the encryption information was not available. With this update, **stratisd** maintains the volume key in its own process keyring. The key is automatically loaded when the pool is unlocked or when the service starts with an existing encrypted pool. To ensure security, the key is removed from the keyring when the **stratisd** process exits or when the pool is stopped or destroyed. If the pool is a V2 encrypted pool and the volume key is not present in the **stratisd** process keyring, **stratis-cli** displays an alert in its pool listing.

Jira:RHEL-125937[1]

snapm rebased to 0.7.0

The **snapm** package has been rebased to upstream version 0.7.0. This version provides important fixes and enhancements, most notably the following:

- The new Mount Manager mounts and unmounts entire snapshots. You can run commands or interactive shells inside mounted snapshot sets by using the **snapset {mount, umount, exec, shell}** subcommands.
- The Difference Engine was added to compare snapshot sets or to compare against the running system. You can specify output formats, such as **paths**, **full**, **short**, **json**, **diff**, **summary**, and **tree**.
- The performance of the Stratis plugin was improved. With this update, the plugin queries the D-Bus every 5 seconds and caches the results internally. This improvement significantly reduces the time to discover Stratis snapshots.

Jira:RHEL-137376[1]

Multipath automatically removes unmapped LUNs

Before this update, multipath devices remained in the system if you did not remove SCSI devices before disconnecting a LUN. This sometimes resulted in queued I/O or incorrect writes if the LUN was repurposed.
With this update, the **purge_disconnected** option is available in the **defaults**, **devices**, and **multipaths** sections of the **multipath.conf** file. When you set this option to **yes**, the **multipathd** daemon automatically removes disconnected SCSI devices from the system.

Jira:RHEL-141287

6.11. HIGH AVAILABILITY AND CLUSTERS

Review new features and enhancements for high availability and clusters in Red Hat Enterprise Linux 10.2.

Ability to add descriptions to cluster resources and elements

Previously, there was no built-in method in **pcs** to add supplemental text descriptions directly to resources and other cluster elements. This limited the ability of administrators to document, provide context, or aid in troubleshooting elements within the Pacemaker cluster.
With this enhancement, a new command, **pcs cib element description**, is available.

As a result, you can add brief text descriptions to a wide range of CIB elements that support the description attribute, including primitive resources, groups, clones, bundles, ACL permissions, ACL roles, alerts, alert recipients, and nodes. For a more intuitive experience, two new aliases are also available: **pcs resource description** and **pcs stonith description**.

Jira:RHEL-7670[1]

Validation added for resource and stonith meta attribute names

Previously, when configuring resource or stonith devices, a user could set meta attributes that were not recognized by the cluster. This led to silent configuration errors where the invalid attributes were accepted without warning but had no effect on cluster resource handling.
With this enhancement, meta attribute names for primitive and stonith resources are validated against the provided cluster meta attributes definition.

As a result, a warning is printed when invalid meta attributes are used with the following commands:

- **pcs resource|stonith create**
- **pcs resource|stonith meta**
- **pcs resource|stonith defaults set create**
- **pcs resource|stonith defaults set update**

Jira:RHEL-7673

Warning added when disabling cluster fencing

Before this update, users could disable the cluster's fencing mechanism by setting the cluster property **stonith-enabled** to false without receiving any warning. This could inadvertently leave the cluster in an unsupported and unsafe state.
With this enhancement, the cluster management utility includes a safety check.

As a result, when you attempt to disable fencing using **stonith-enabled=false** the utility displays a warning message informing you that the cluster fencing mechanism will be lost.

Jira:RHEL-84120

The **portblock** resource agent now supports**nftables**

Previously, the **portblock** resource agent relied on **iptables** for managing port access. Since **iptables** is now primarily a wrapper for **nftables** and is slated for removal in future releases, a transition to native **nftables** support was necessary.
With this enhancement, the **portblock** resource agent now supports **nftables** natively.

As a result, **nftables** is used by default for port blocking operations. For environments that still require the legacy behavior, you can manually switch back to **iptables** by setting the firewall resource parameter to **iptables**.

Jira:RHEL-116152

6.12. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Review new features and enhancements for dynamic programming languages, web and database servers in Red Hat Enterprise Linux 10.2.

MariaDB 11.8 was added

MariaDB 11.8 packages are available in RHEL 10.2.
Notable changes over the previously available version 10.11 include:

- By default, MariaDB 11.8 uses the **utf8mb4** character set instead of **latin1** and legacy **utf8** to ensure full Unicode support.
- Vector support was added to support machine learning. This includes the **VECTOR(N)** data type and the following functions:
  - **VEC_DISTANCE()**
  - **VEC_DISTANCE_EUCLIDEAN()**
  - **VEC_DISTANCE_COSINE()**
  - **Vec_FromText(*json_array*)**
  - **Vec_ToText(*vector_column*)**
- The **mariadb-dump** and **mariadb-import** utilities natively support parallel operations. Specify the **--dir** and **--parallel** options to dump or load multiple databases simultaneously.
- The upper limit of the **TIMESTAMP** data type was increased from **2038-01-19** to **2106-02-07** while still using 4 bytes of storage.
- The **UUID_v4()** and **UUID_v7()** functions were added.
- The JSON handling was improved. This includes new functions, such as **JSON_SCHEMA_VALID()**.
- The following system variables were added to define the maximum storage for temporary tables and other internally created temporary files:
  - **max_tmp_session_space_usage** limits the disk space used per session
  - **max_tmp_total_space_usage** limits the total disk space used by the MariaDB server instance
- The **des_encrypt** and **des_decrypt** configuration file parameters are deprecated and will be removed in a future MariaDB release.

Notable breaking differences:

- The following utilities were renamed but symbolic links were created for backward compatibility:
  - **mysql** > **mariadb**
  - **mysqldump** > **mariadb-dump**
  - **mysqladmin** > **mariadb-admin**

  If you still use the previous names of these utilities, they display deprecation warnings.
- The **innodb_defragment** configuration parameter is no longer supported. Remove it from your configuration files.

For more information about MariaDB, see Using MariaDB.

To install the new packages, enter:

```
# dnf install mariadb11.8-server
```

If you want to upgrade from MariaDB 10.11, see Upgrading from a RHEL 9 version of MariaDB 10.11 to MariaDB 11.8.

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

Jira:RHEL-115468[1]

PostgreSQL 18 was added

PostgreSQL 18 packages are available.
Notable changes:

- The new Asynchronous I/O (AIO) subsystem provides up to three times faster data reads. You can enable this subsystem by setting the **io_method** variable.
- The MD5 authentication method is deprecated and will be removed in a future major PostgreSQL release.
- By default, data page checksums are enabled in PostgreSQL 18. If you upgrade from a previous version with data page checksums disabled, you must either enable the feature before the update or disable it during the upgrade. For further details, see Upgrading from a RHEL 9 version of PostgreSQL 16 to PostgreSQL 18.
- PostgreSQL 18 supports native OAUth 2.0 single sign-on authentication.
- The database service supports Federal Information Processing Standards (FIPS) mode validation for regulated environments.
- The **pg_upgrade** utility preserves statistics during major release upgrades and significantly faster reaches full performance after an upgrade.

Jira:RHEL-116546[1]

New **ruby4.0** runtime with database connectors

The **ruby4.0** runtime provides a Ruby 4.0 stack with database connector support. As a result, you can select the **ruby4.0** runtime alongside the existing Ruby stack to develop and run Ruby applications with supported database connectivity.

Jira:RHEL-133550[1]

New Python 3.14 stack is available

The **python3.14** stack with essential packages is available in RHEL 10.2. This new alternative stack provides Python 3.14 to develop and run applications while staying on the RHEL 10 minor release.

Jira:RHEL-120788[1]

6.13. COMPILERS AND DEVELOPMENT TOOLS

Review new features and enhancements for compilers and development tools in Red Hat Enterprise Linux 10.2.

glibc memstream documentation updated for SEEK_END behavior

The **glibc** memstream documentation explains how **open_memstream** handles seeking and the current position when using **SEEK_END**. The updated text clarifies how writing at the end of the buffer behaves, in line with the implementation-defined behavior introduced in POSIX Issue 8.

Jira:RHEL-65838

New OpenTelemetry PMDA enables OTLP metric ingestion into PCP

A new Performance Metrics Domain Agent, **pmdaopentelemetry**, is available to ingest OpenTelemetry metrics into Performance Co-Pilot (PCP). This enhancement bridges the OpenTelemetry ecosystem with PCP by dynamically creating PCP metrics from configured endpoints that export data in OTLP JSON format. The PMDA replaces the legacy **pmdajson** for OpenTelemetry use cases.

Features include:

- Support for HTTP/HTTPS endpoints, local files, and executable script data sources.
- Dynamic metric namespace with automatic addition and removal of metrics without requiring a restart.
- Regex-based INCLUDE, EXCLUDE, and OPTIONAL rules for filtering metrics and labels.
- Configurable metadata overrides for PCP types, semantics, and units.
- Support for OpenTelemetry metric types such as sum, gauge, histogram, and summary.
- Concurrent multi-source data collection using a thread pool.
- Persistent metric and instance identifiers across restarts.
- Built-in control metrics for per-source monitoring and diagnostics.

Jira:RHEL-83866

PCP REST API supports exporting metrics in OpenTelemetry JSON format

The **pmproxy** service supports exporting Performance Co-Pilot (PCP) metrics in OpenTelemetry JSON format through the existing **/metrics** REST API endpoint. When a client includes the **Accept: application/json** header in the HTTP request, **pmproxy** returns metrics in the OpenTelemetry **resourceMetrics** JSON structure instead of the default OpenMetrics text format.
This enhancement enables direct integration of PCP metrics with OpenTelemetry-based monitoring solutions without requiring additional format conversion. The existing OpenMetrics text format remains the default when the **Accept: application/json** header is not specified. Features include:

- Support for exporting PCP metrics in OpenTelemetry **resourceMetrics** JSON format through the **/metrics** endpoint.
- Content negotiation using the **Accept: application/json** HTTP header.
- Automatic conversion of PCP metric semantics, types, and labels to OpenTelemetry equivalents.
- Conversion of PCP units to the Unified Code for Units of Measure (UCUM) format.
- Compatibility with existing OpenMetrics text format as the default response format.

Jira:RHEL-85456

New tool pcp2opentelemetry introduces OpenTelemetry data export

With this update, a new tool, **pcp2opentelemetry**, is introduced for exporting both real-time and archived Performance Co-Pilot (PCP) data in the OpenTelemetry format. This tool extends OpenTelemetry support within PCP, similar to **pcp2openmetrics**, and is part of the ongoing support for OpenTelemetry in PCP v7. By using this tool, you can export PCP data in the OpenTelemetry format. It boosts compatibility with other tools within the OpenTelemetry ecosystem and offers a more adaptable and integrated method for managing performance data.
Jira:RHEL-85457

New PMDA for SAP HANA database metrics

With the Performance Co-Pilot (PCP), a new Performance Metrics Domain Agent (PMDA) is available for monitoring SAP HANA databases. You can now use PCP to collect and analyze metrics from SAP HANA, enabling improved visibility into database performance and behavior. This enhancement helps administrators monitor the SAP HANA workloads by using standard PCP tools and workflows.
Jira:RHEL-85725

Rebase llvm toolset to version 21

The **llvm** toolset has been rebased to version 21 in RHEL 10.2. This rebase provides updated compiler and tooling features for building and optimizing applications that depend on **llvm**.
As part of this change, dependent packages in RHEL 10 have been rebuilt against **llvm** 21 to ensure compatibility with the updated toolset.

The notable changes are:

- The **nocapture** function attribute is replaced by the more expressive **captures(none)** attribute in LLVM IR, clarifying pointer capture semantics.
- Constant expression forms of several arithmetic instructions, including **mul**, are removed in favor of using regular instructions, simplifying IR and optimizations.

- Inline assembly calls no longer accept **label** operands. The **callbr** instruction must be used instead, which clarifies semantics for indirect labels.
- New **fmaximum** and **fminimum** operations are supported in the **atomicrmw** instruction, aligning atomic floating-point operations with **llvm.maximum** and **llvm.minimum** behavior.
- Multiple back ends, including AArch64, AMDGPU, RISC-V, PowerPC, and others, receive code generation improvements, new ISA extensions, and bug fixes that can result in better performance and broader hardware support.

Jira:RHEL-100887

PCP supports PUSH model for pmlogger

PCP supports a push model for **pmlogger** that enables remote archival of performance metrics data by using an HTTP REST API. Previously, centralized logging required administrators to reconfigure the central system to pull data from each newly added host. With the push model, each host streams archived data directly to a centralized **pmproxy** server in real time. This approach simplifies scaling and removes the need to store archives locally on remote systems. Additional key features include the following:

- Centralized Storage: Configure multiple remote systems identically to store all archives on a single **pmproxy** server, simplifying configuration, management, and backup.
- No Local Storage Required: Remote hosts log metrics without requiring local disk space for archives.
- Real-Time Streaming: Archive data is transmitted immediately as it is collected, enabling near real-time analysis.
- Network Resilience: Built-in error handling and retry mechanisms during network interruptions.

Jira:RHEL-104669[1]

Enhanced gcov function coverage summaries in gcc

Before this update, **gcov** function summaries only reported the number of lines executed and did not include details about branch or call coverage within the function.
With this enhancement, requesting function summaries using the **-f** option now includes data on branches taken and function calls made within the profiled function. This provides a more comprehensive view of function-level test coverage.

Jira:RHEL-105464[1]

glibc fortification support 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 detect some buffer errors before running the program.
With this update, attribute access annotations is added to **inet_ntop** and **inet_pton**, enabling the compiler to warn about potential buffer misuse at compile time. As a result, these APIs are now covered by Source Fortification, which improves their security and reliability.

Jira:RHEL-111115[1]

Rust Toolset is rebased to versions 1.92.0

RHEL 10.2 rebases the **rust-toolset** Application Stream to version 1.92.0, providing an updated Rust compiler and associated tooling for developing and running Rust applications. This rebase continues the rolling Application Stream model, where only the latest **rust-toolset** version is supported.
Notable enhancements include:

- Reliable debugging through default emission of unwind tables on Linux, even when compiling with **-Cpanic=abort**, which enables more accurate backtraces.
- Expanded systems programming support, including full **i128** and **u128** support in extern "C" functions and the ability to create raw pointers to union fields using **&raw** in safe code.
- Enhanced code safety with the new **dangling_pointers_from_locals** lint, which warns about returning dangling raw pointers derived from local variables.
- Improved code clarity with the **mismatched_lifetime_syntaxes** lint, which highlights potentially confusing lifetime relationships that are hidden by lifetime elision rules.
- Workflow improvements in Cargo, which supports workspace-level publishing with **cargo publish --workspace** and automatically handles dependency ordering for multi-crate projects.
  Rust Toolset is delivered as a rolling Application Stream, and only the latest rust-toolset version is supported. For more information about Rust Toolset life cycle and support, see the Red Hat Enterprise Linux Application Streams Life Cycle .

Jira:RHEL-111845

glibc **freopen** behavior fixes and test improvements

The **freopen** function behaves more reliably and correctly under various usage scenarios. The function no longer leaks memory on failure, preserves and applies file descriptor flags such as **O_CLOEXEC** correctly, and sets the stream orientation properly when a character set is specified.

Jira:RHEL-115823[1]

Improved vectorized math performance in**glibc**

The **glibc** vectorized math library (**libmvec**) was upgraded to the upstream 2.40 version. The upstream **glibc** 2.40 release adds 55 additional vectorized math routines that were previously not available in the RHEL **glibc** based on version 2.39.
As a result, vectorized code compiled with the **-ffast-math** build option on AArch64 now benefits from these functions and might use symbols with the **glibc** version 2.40.

Jira:RHEL-118273[1]

Boost URL library available through boost in CRB

The **boost-url** shared library is available as part of the main **boost** package in the CodeReady Builder (CRB) repository. This change resolves the missing **boost-url** subpackage that blocked some dependent builds in earlier releases.
Because **boost-url** is a dependency of the **boost** metapackage, it is shipped with **boost** instead of as a separate repository entry. In RHEL 10.2, the **boost-1.83.0-7.el10** build ensures that **boost-url** is included in the product listing, and installing **boost-devel** also provides the headers and libraries needed to build applications that rely on Boost.URL.

Jira:RHEL-124169

Performance Co-Pilot 7.0.3 enhancements

pcp-7.0.3-1.el10 in RHEL 10.2 introduces enhancements to monitoring capabilities, including new metric sources and improved sample resolution.
The update adds new Performance Metrics Domain Agents (PMDAs), expands **dstat** plugin coverage, and improves timestamp granularity for collected samples. The following enhancements are included in this update:

- Added a new PMDA to collect SAP HANA database metrics.
- Added a new PMDA for RoCE devices.
- Added a new PMDA to collect OpenTelemetry metrics and a native PCP to OpenTelemetry exporter utility.
- Enhanced the **pmproxy** REST API with a **pmlogger** push mode to send data to a centralized location. This simplifies configuration by allowing you to use the same setup on multiple machines that collect performance data.
- Added new **dstat** plugins to monitor AMD and NVIDIA GPUs.
- Added a new **dstat** plugin to monitor vLLM cache utilization.
- Enabled nanosecond timestamp resolution on collected samples.

Jira:RHEL-124897

The Red Hat Build of OpenJDK 25 integrates with the **crypto-policies** package for secure system property handling

With this update, the Red Hat Build of OpenJDK 25 for RHEL integrates with the RHEL **crypto-policies** package. This enhancement ensures secure system property handling and improves the security of Java applications running on RHEL by loading additional configuration files based on Red Hat system properties. This change also adds FIPS support using NSS.

Jira:RHEL-128409[1]

glibc updates the **bg_BG** locale for Bulgaria's adoption of the euro

The **glibc** package uses the euro currency symbol for the **bg_BG** locale to reflect Bulgaria's adoption of the euro as of 1 January 2026.
As a result, applications that use the **bg_BG** locale display currency values with the updated euro symbol.

Jira:RHEL-137184

Croatia locale uses the euro currency symbol in **glibc**

The **glibc** package now uses the euro currency symbol for the **hr_HR** locale in RHEL. This change aligns Croatian locale data with the country's current official currency.
As a result, applications that rely on **glibc** locale information for the **hr_HR** locale now display the up-to-date euro currency symbol instead of the former Croatian kuna.

Jira:RHEL-140103

Improved **trylock** performance in **glibc** for heavily contended multi-core workloads

With this enhancement, the **glibc** package optimizes the **trylock** implementation for workloads with high thread counts on multi-core systems, improving **trylock** throughput under heavy contention.
Jira:RHEL-139419

glibc adds **RTLD_DI_ORIGIN_PATH** to prevent buffer overflows

The **RTLD_DI_ORIGIN_PATH dlinfo** request type in **glibc** accepts the size of the destination buffer when retrieving the shared object origin path. This request type helps avoid buffer overflows when obtaining the shared object origin path.
The behavior of the existing **RTLD_DI_ORIGIN** request type remains unchanged.

Jira:RHEL-146428[1]

6.14. IDENTITY MANAGEMENT

Review new features and enhancements for Identity Management (IdM) in Red Hat Enterprise Linux 10.2.

The entry and DN cache auto-sizing considers the number and size of DS databases

With this update, the auto-sizing feature for entry and DN caches adapts its tuning when a Directory Server instance uses multiple databases of different sizes. The cache size matches the database size, allocating more physical resources to larger databases.
Jira:RHEL-18041

You can pin large groups in the Directory Server entry cache to prevent groups eviction

With this update, Directory Server introduces a new configuration attribute, **nsslapd-cache-pinned-entries**, in backend configuration entries to pin the largest groups in the entry cache. You can set the number of entries that you want to pin by using the **nsslapd-cache-pinned-entries** attribute. These group entries are only evicted when modifying the group or when bringing the backend down. The default value is **0** meaning no group entries are pinned.
Jira:RHEL-58682

The **ipa-client-automount** utility now supports the **--domain** option

Before this update, the **ipa-client-automount** utility relied on the host's current DNS domain for service discovery. This caused issues in cross-domain environments where the client host resided in a different DNS domain than the Identity Management (IdM) server, often requiring administrators to manually configure numerous server settings in multiple locations.
With this update, **ipa-client-automount** introduces the **--domain** option. This allows users to explicitly define the IdM domain to be used for DNS discovery during the automount configuration.

As a result, installation efficiency and reliability are improved for complex network topologies.

Jira:RHEL-86030[1]

Directory Server supports online TLS certificate refresh without the server restart

With this update, you can update server certificates on a running instance and trigger a certificate refresh without stopping the **dirsrv** service. After deploying new certificates, you can use the **dsconf *<instance_name>* config refresh-certs** command to activate them for new incoming TLS connections, enabling smoother, more automated certificate renewal processes with less downtime. Existing LDAP connections are not explicitly closed. However, if the CA certificate has changed, some existing LDAPS connections might be terminated by clients with **SERVER_DOWN** errors. This occurs when the clients expect the previous certificate while the server renegotiates encryption with the new one.
Jira:RHEL-86320

Directory Server supports dynamic groups

With this update, you can define group membership based on LDAP search filters, similar to OpenLDAP, instead of managing static member lists. Using search filters to define group membership provides more flexible and scalable access control. Membership is automatically calculated from LDAP search URLs when you configure a dedicated object class, URL attribute, and list attribute. As a result, Directory Server introduces the following configuration attributes under **cn=config,cn=ldbm database,cn=plugins,cn=config**:

- **nsslapd-dynamic-lists-enabled** enables dynamic groups. Defaults to **off**.
- **nsslapd-dynamic-lists-oc** defines which object classes an entry must have to be considered a dynamic entry. Defaults to **groupOfURLs**.
- **nsslapd-dynamic-lists-url-attr** specifies which attribute inside a group entry contains the LDAP URL used to define the dynamic membership. Defaults to **memberUrl**. The attribute can be multi-valued if your schema allows it.
- **nsslapd-dynamic-lists-attr** specifies the attribute that contains the distinguished name (DN) of the entries that match the LDAP URI set in **nsslapd-dynamic-lists-url-attr**. Defaults to **member**.

Jira:RHEL-86534

Delete all replication conflict entries at once with**dsconf**

With this update, you can use the **dsconf *<instance_name>* repl-conflict delete-all "<suffix_name>"** command to delete all replication conflicts in bulk. Before this update, each conflict had to be deleted individually by using **dsconf *<instance_name>* repl-conflict delete**. Now, you can delete all replication conflicts in a single operation by using **dsconf**.
Alternatively, you can try to resolve conflicts instead of deleting them. For details, see Solving common replication problems.

Jira:RHEL-99331[1]

Directory Server validates LDIF files before erasing backend data during import

Before this update, importing the wrong LDIF file would erase the backend first and only report errors after processing the entire file, potentially causing data loss.
With this update, the server performs early validation when importing LDIF files to detect mismatched or incorrect files before erasing the existing backend database. If the LDIF does not contain the expected suffix entry, the import terminates immediately with a clear error message, leaving the existing backend data intact.

Jira:RHEL-106849

dsctl dbverify provides clearer output when a specified backend does not exist

The **dsctl dbverify** command, used to verify the integrity of a Directory Server database, provides explicit feedback depending on the database backend type. For Lightning Memory-Mapped Database (LMDB) backends, the command displays a warning that the verification is always reported as successful because LMDB has built-in integrity protection. As a result, administrators can distinguish between a missing backend and a genuinely successful verification when running **dsctl dbverify**.
Jira:RHEL-107003

MemberOf plugin supports scoping for specific groups

With this update, you can configure the MemberOf plugin to monitor only selected groups for membership evaluation. Previously, MemberOf plugin processing was controlled at the suffix level, which included all groups under a configured suffix. By defining a group scope, you can target list of groups or create exceptions for specific groups. This improves performance by avoiding unnecessary plugin operations on irrelevant entries.
MemberOf plugin introduces the following **multi-valued** configuration attributes under **cn=MemberOf Plugin,cn=plugins,cn=config**:

- **memberOfSpecificGroupFilter** sets an LDAP search filter to select the group entries the plugin should process.
- **memberOfExcludeSpecificGroupFilter** sets an LDAP search filter to select the group entries to be excluded from plugin processing.
- **memberOfSpecificGroupOC** sets the object class of the group entries the plugin should process.

Jira:RHEL-109113[1]

Directory Server supports post-quantum cryptography (PQC) keys

With this update, Directory Server supports TLS certificates that use **ML-DSA-44**, **ML-DSA-65**, and **ML-DSA-87** keys. This enables adoption of post-quantum cryptography standards to help protect your directory against potential quantum computing attacks.
Jira:RHEL-110192

You can configure external password reset agents in IdM

When integrating Identity Management (IdM) with a third-party application that does not support Kerberos authentication, you can define a dedicated system account for the application to securely reset user passwords. Notably, these resets do not trigger the "password change required" flag, ensuring a seamless login experience for the end user. The system account authenticates by using LDAP.
As a result, organizations can integrate their own secure password management solutions directly with IdM.

Jira:RHEL-110204

You can specify an IdM server from which to update the local CA trust store

With this update, the **ipa-certupdate** tool includes a new **--force-server *<server_fqdn>*** option. Before this update, an Identity Management (IdM) client only connected to its default IdM server, specified in the **/etc/ipa/default.conf** file, when updating the local CA trust store. If this default server was down or unreachable, the **ipa-certupdate** command failed. As a result, administrators can ensure successful trust store updates and maintain service continuity, even if the primary server is unavailable.
Jira:RHEL-113778

samba rebased to 4.23.0

The **samba** packages, which provide file and print services using the SMB protocol, have been rebased to upstream version 4.23.0. This version provides important fixes and enhancements, most notably the following:

- SMB3 UNIX Extensions are enabled by default to provide support for POSIX semantics, such as proper POSIX permissions and symlink handling, for UNIX and Linux clients.
- Experimental support for SMB3 connections over Quick UDP Internet Connections (QUIC) is introduced. Configurable through **client smb transports** and **server smb transports**, this allows for secure SMB traffic over UDP port 443, which is ideal for remote access.
- The new **smb_prometheus_endpoint** utility exports Samba server metrics in a Prometheus-compatible format to facilitate performance and status monitoring.
- The **samba-tool domain backup --no-secrets** command explicitly removes confidential attributes, such as BitLocker recovery data and KDS root keys, from backups.
  For a complete list of changes, see Samba 4.23.0 Available for Download.

Jira:RHEL-114545

IdM password policies support libpwquality character credit options

Identity Management (IdM) password policies support four new options (**--dcredit**, **--ucredit**, **--lcredit**, and **--ocredit**) based on the **libpwquality** credit system. A negative value sets the minimum number of characters of that type required in a password; a positive value provides a credit toward the minimum password length. These options are mutually exclusive with **--minclasses** and offer a more granular way to enforce per-class character requirements. As a result, administrators can configure specific character type minimums in IdM password policies, for example, to satisfy DISA STIG compliance requirements.
For more information, see Additional password policy options in IdM .

Jira:RHEL-119481[1]

ipa rebased to 4.13.0

The **ipa** packages have been rebased to upstream version 4.13.0. This version provides important fixes and enhancements, most notably the following:

- A new responsive and intuitive beta interface is available as a Technology Preview. You can experiment with it and provide feedback.
- You can use the **ipa-idrange-fix** tool to identify users and groups outside current ID ranges and propose new ranges to include them.

- The requirement for unique Certificate Authority (CA) subject names is relaxed, which enables duplicates under specific trust and nickname conditions.
- Random serial numbers (RSNv3) are enabled by default, and the system automatically removes certificates 30 days after they expire.
- To modernize new deployments, Network Information Service (NIS) server emulation in Identity Management (IdM) is removed. Note that NIS client support was removed in Red Hat Enterprise Linux (RHEL) 9. Additionally, the Schema Compatibility Tree plugin is deprecated and might be removed in a future major release.
- The platform supports the full 32-bit ID range space.
- This release resolves over 170 bugs and improves overall system performance and stability.

Jira:RHEL-120956[1]

nsslapd-haproxy-trusted-ip now supports CIDR notation

With this update, you can use Classless Inter-Domain Routing (CIDR) notation to define ranges of trusted IP addresses instead of manually listing each address. You can now specify multiple CIDR ranges, as well as a mix of individual IPs and ranges. Example multi-valued configuration:

```
nsslapd-haproxy-trusted-ip: 2001:db8::/32
nsslapd-haproxy-trusted-ip: 192.168.1.0/24
nsslapd-haproxy-trusted-ip: 192.168.2.50
```

RHEL-121208

Jira:RHEL-121208

cepces rebased to 0.3.12

The **cepces** package, which provides a certificate enrollment client for Microsoft Active Directory Certificate Services (AD CS), has been rebased to upstream version 0.3.12. This version provides important fixes and enhancements, most notably the following:

- Support for GSSAPI channel bindings to bind Kerberos authentication to the TLS (HTTPS) tunnel is available. This is required for compatibility with Windows Server 2025, which enforces stricter security requirements for SOAP-based certificate enrollment web services (CEP/CES) by default.
- Authentication handshake failures when connecting to modern Windows environments that have TLS channel binding and Kerberos security policies enabled are fixed.
- Updates to the **cepces-submit** helper ensure smoother communication with the **certmonger** service during automated certificate renewal cycles.

Jira:RHEL-121729

Support for generating LWCA certificates and private keys on an HSM

For installations using a hardware security module (HSM), Lightweight CA (LWCA) certificates and private keys are now generated on the HSM. This provides the same hardware-level security for the private keys as the root CA private key. The LWCA private key is generated on the HSM with the HSM token name as the prefix, for example **mytoken:lwca**.

Jira:RHEL-126761

Automated services no longer reset account lockout counters

This update ensures that automated services like **crond** and **systemd-user** are prevented from unlocking accounts locked by **faillock**. Previously, these services would automatically clear the "failed login" counter when they ran, which could allow a malicious actor to keep guessing passwords without being permanently locked out. With this release, once an account is locked by a security policy, it remains locked until the timeout expires or an administrator intervenes, regardless of any background system activity.

Jira:RHEL-130871[1]

ansible-freeipa rebased to 1.16.0

The **ansible-freeipa** packages, which provide Ansible modules and roles for Identity Management (IdM), have been rebased to upstream version 1.16.0. This version provides important fixes and enhancements, most notably the following:

The **sysaccount** module (**ipasysaccount**) creates and manages system accounts in IdM. The **role** module (**iparole**) supports system accounts as role members, so you can assign privileges such as user password management to those accounts in playbooks. You can, for example, use system accounts to integrate IdM with an external password reset management solution. For more information, refer to the **sysaccount** and **role** module READMEs.

The **ipapasskeyconfig** module is available in the **ansible-freeipa** collection. You can use this module to configure whether passkey authentication in IdM requires user verification, such as a PIN, when users authenticate with a passkey device. Additionally, the **ipauser** module supports **passkey** as a user authentication type, and the **ipaservice** and **ipahost** modules support **passkey** as an authentication indicator.

Jira:RHEL-139147

ansible-freeipa adds support for thepasskey authentication type in management modules

With this update, the **ipaconfig**, **ipahost**, **ipaservice**, and **ipauser** modules support the **passkey** authentication type for IdM resources. This enables you to manage Passkey device authentication directly through your Ansible playbooks by setting the authentication type to **passkey**.

Jira:RHEL-139258

389-ds-base rebased to 3.2.0

The **389-ds-base** package, which provides an enterprise-class LDAP server, has been rebased to upstream version 3.2.0.

Jira:RHEL-139826

The Certificate System now supports ML-DSA keys and signatures

You can now install a Certificate System (CS) that uses Module-Lattice-based Digital Signature Algorithm (ML-DSA) for both key types and signatures. Because ML-DSA is standardized by NIST to withstand future quantum computing threats, the CS can now generate and manage quantum-resistant certificates. This release supports ML-DSA at three NIST-defined security levels: ML-DSA-44, 65, and 87.

Jira:RHEL-143038

pki rebased to 11.9

The **pki** packages have been rebased to upstream version 11.9. This version provides important fixes and enhancements, most notably the following:

- Support for ML-DSA (Module-Lattice-based Digital Signature Algorithm) profiles is available. This enables the PKI to issue and manage certificates using post-quantum cryptographic algorithms, preparing the environment for future security requirements.
- The Jackson JSON processing libraries are updated to improve performance and security during metadata serialization.
- The **pki-server** and associated CLI tools include stability updates to better handle service state transitions and improve the reliability of trust store synchronization in complex topologies.
- A race condition that caused **ipa ca-add** to fail with a "500 Internal Server Error" when adding multiple Sub-CAs in rapid succession is resolved. With this update, the CA engine correctly synchronizes authority initialization with signing certificate availability, which prevents API timeouts during high-volume operations.
- A regression where enabling the **nuxwdog** watchdog prevented the PKI service from starting is fixed. The **pki-server-nuxwdog** utility correctly interfaces with **systemd-ask-password**, enabling users to provide required credentials at startup when a password file is missing.
- An issue where the PKI server failed to issue certificates when a Sub-CA was specified is resolved. This fix ensures the certificate request pipeline correctly identifies and utilizes Sub-CA signing keys, which restores full functionality to multi-tier CA environments.

Jira:RHELDOCS-21885[1]

6.15. SSSD

Review new features and enhancements for SSSD in Red Hat Enterprise Linux 10.2.

Recursive deletion for computer objects added to adcli

The **adcli delete-computer** command supports the **--recursive** option to delete computer objects from Active Directory, including their child objects. Previously, attempting to delete a computer object that contained child objects, such as metadata for BitLocker drive recovery, failed with a **CANT_ON_NON_LEAF** error in AD. With this update, users can cleanly delete computer objects that contain child objects using **adcli**.

Jira:RHEL-16141

sudo rebased to sudo-1.9.17p2

The **sudo** packages have been rebased to upstream version 1.9.17p2, which includes the following notable bug fixes and enhancements:

- The **sudoers** file supports regular expressions.
- The **log_subcmds** and **intercept** options are supported.
- The **json_compact** logging is supported.
- Privilege listing is enhanced.

- Added the **cmddenial_message sudoers** option.
- The **sudoers** LDAP schema now allows **sudoUser**, **sudoRunasUser**, and **sudoRunasGroup** to include UTF-8 characters.
- Added a new **-N** (no-update) command-line option to **sudo**.
- The following **sudoers** settings can be used to support more fine-grained I/O logging:
  - **log_stdin**
  - **log_stdout**
  - **log_stderr**
  - **log_ttyin**
  - **log_ttyout**

Jira:RHEL-112100

6.16. DESKTOP

Review new features and enhancements for desktop in Red Hat Enterprise Linux 10.2.

The display time for login error messages is extended

Before this update, some short error messages on the login screen disappeared too quickly to be read. As a consequence, users missed important login feedback. With this update, the display time for short error messages is extended. As a result, these messages remain visible for a longer period of time.
Jira:RHEL-11918

papers rebased to 48.4

The **papers** document viewer is rebased to version 48.4. This version provides important fixes and enhancements, most notably the following:

- Support for the **libspelling** library is added.
- Stability of **papers** is improved.
- Various UI improvements are in place.
- Support for the PostScript and XPS document formats is removed.
- The bookmarks sidebar is removed.
- Translations are updated.

Jira:RHEL-86193

fwupd package is rebased to 2.0.19

The **fwupd** package, which updates firmware on your system, has been rebased to upstream version 2.0.19. This version provides important fixes and enhancements, most notably the following:

- Applied important fixes to the various firmware loaders, such as PE/COFF and MTD.
- Fixed issues affecting Intel GPUs and docks from USI, Lenovo, Dell, and HP.
- Improved deployments for UEFI KEK, db, and dbx updates.
- Added support for client-side phased update deployment, post-quantum cryptography (PQC) signatures, and additional devices including NVIDIA ConnectX, Jabra Evolve2, Framework QMK, Copilot devices, Huddly C1, and SteelSeries Arctis Nova.

Jira:RHEL-110760[1]

libinput rebased to version 1.30

The **libinput** package is rebased to upstream version 1.30. This version provides important fixes and enhancements, most notably the following:

- 3-finger dragging for touchpads and a new sticky drag-lock feature for the tap-and-drag setting are supported.
- Configuration for mapping an eraser button for tablet tools with a hardcoded eraser button is supported. This behavior applies to most Microsoft-compatible tablets other than Wacom tablets.
- Tablets without physical LEDs to indicate the tablet pad mode are handled correctly.
- Configuration of the accessible tablet area on external tablets is supported. As a result, you can reduce the available physical area to better match it to the intended use case.
- Many device-specific updates are added to accommodate custom behavior required by specific devices.
  For any new configuration option, **libinput** provides the option, but it must be set by the respective compositor. Depending on the compositor, some configuration options might not be available directly to the user.

Jira:RHEL-136390

Flatpaks are the default delivery method for Mozilla Firefox and Thunderbird

With this update, the default delivery method for Mozilla Firefox and Thunderbird is changed from RPM packages to Flatpaks. Anaconda, the RHEL installer, preinstalls these Flatpaks by default.
If your system is subscribed to Red Hat, you do not need to provide your Red Hat credentials when accessing the Red Hat Flatpak Registry. If you use the Red Hat Flatpak Registry on an unsubscribed system, follow the official guidelines.

Because there might be use cases where Flatpaks do not fit well, Red Hat will continue to provide and support **firefox** and **thunderbird** RPM packages in the **AppStream** repository for the lifetime of RHEL 10. If you identify any of these use cases, contact Red Hat Support. Alternatively, you can provide your feedback in the RHEL-160615 Jira ticket.

You can change the delivery method in Anaconda from Flatpaks back to RPM packages by following the process outlined in the documentation. For example, use the following configuration to preinstall the **firefox** RPM package instead of the Flatpak:

```
%packages
@^graphical-server-environment
-redhat-flatpak-preinstall-firefox
firefox
%end
```

Jira:RHEL-139533

6.17. THE WEB CONSOLE

Review new features and enhancements for the web console in Red Hat Enterprise Linux 10.2.

cockpit rebased to version 356

The **cockpit** packages have been rebased to version 356, which provides many improvements and fixes compared to version 344 in RHEL 10.1, most notably:

- Timers created by the RHEL web console are executed directly by the **/bin/sh** system shell, and you can edit them.
- The health dashboard shows a warning if the last shutdown or reboot was unclean.
- You can override the RHEL web console branding with a custom configuration in the **/etc/cockpit/branding.css** file.
- Support for the **pam_cockpit_cert** PAM module in the **/etc/pam.d/cockpit** file, which is redundant since version 248, is removed. If you still use the module in your configuration, you must remove it manually.
- The web console lists additional ports in a firewall zone, each in its own row, and you can delete them individually.
- Support for TLS is removed from the **cockpit-ws** subpackage. Instead, containers run the **cockpit-tls** program and directly connect to the **cockpit-ws** server.
- You can detach the VNC console viewer of a virtual machine into its own window.
- The web console no longer adds both SPICE and VNC graphics when creating new virtual machines, but only VNC.
- You can shut down and restart virtual machines with a single action from the web console.
- The **cockpit-podman** plug-in supports the quadlet lifecycle and shows inactive quadlets.
- You can create empty files in the web console file manager.

Jira:RHEL-112867

6.18. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Review new features and enhancements for Red Hat Enterprise Linux system roles in Red Hat Enterprise Linux 10.2.

The ha_cluster RHEL System Role now exports cluster constraint variables

Previously, the **ha_cluster** RHEL System Role did not include detailed constraint information in its exported data.
With this enhancement, the **ha_cluster** role now includes variables for location, colocation, order, and ticket constraints.

As a result, the following variables are now available in the module output, facilitating better configuration management and role-based automation:

- **ha_cluster_constraints_location**
- **ha_cluster_constraints_colocation**
- **ha_cluster_constraints_order**
- **ha_cluster_constraints_ticket**

Jira:RHEL-46226

The **ha_cluster** RHEL System Role now exports additional cluster configuration variables

Previously, the **ha_cluster** RHEL System Role provided limited visibility into the current cluster configuration.
With this update, the **ha_cluster** role has been expanded to include cluster properties and resource defaults.

As a result, the following variables are now exported, allowing for easier auditing and configuration mirroring:

- **ha_cluster_cluster_properties**
- **ha_cluster_resource_defaults**
- **ha_cluster_resource_operation_defaults**

Jira:RHEL-46227

Disk partition management available to the storage role

With this update, you can manage disk partitions by using the storage role, streamlining storage management. With this unified approach you can add, remove, resize, and format partitions, ensuring consistent and repeatable results.

Jira:RHEL-66738[1]

Support for bootable snapshots with **snapm**

With this update, you can create bootable snapshot sets on platforms that support **snapm**, such as RHEL 9.6 and Fedora 41 or later. You can now set a **bootable** flag when requesting snapshots and boot the system directly from a snapshot.
Jira:RHEL-104931

The **firewall** RHEL system role supports IPv6 addresses within the **ipset_entries**

With this enhancement, you can now use IPv6 addresses within the **ipset_entries** variable when utilizing **hash:ip** or **hash:net** types in playbooks that use the **firewall** RHEL system role. You can

also specify additional **<key>:<value>** pairs of options for **ipset** by using the **ipset_options** variable.
pairs
Due to a limitation of the underlying **firewalld** implementation, you cannot mix IPv4, IPv6, and MAC addresses in the same **ipset_entries** list.

Jira:RHEL-114467[1]

The **sshd** system role supports the**CanonicalMatchUser** option

To provide more granular control over conditional configurations, the **sshd** system role supports the **sshd_CanonicalMatchUser** variable. You can specify whether to evaluate OpenSSH **Match** blocks against a user's initial login name or their final canonical username after the server rewrites it.
As a result, you can consistently apply security policies in environments where external identity providers or local configuration rules modify usernames. This ensures that **Match** blocks accurately reflect the user's identity once the server determines the final canonical username.

Jira:RHEL-127971

Support added for the**fencing-watchdog-timeout** cluster property

Before this update, the high-availability stack primarily supported the **stonith-watchdog-timeout** property for managing watchdog-based fencing. However, future Pacemaker versions replace this property with **fencing-watchdog-timeout**.
With this update, the role handles both the legacy and new property names consistently.

As a result, the role supports future Pacemaker versions and ensures that watchdog-related cluster properties remain functional regardless of which property name you use. The role preserves both **stonith-watchdog-timeout** and **fencing-watchdog-timeout** when creating or pushing CIB configurations.

Jira:RHEL-136597

The **metrics** RHEL system role supports configuring TLS-encrypted connections

With this enhancement, you can use the **metrics** RHEL system role to configure TLS-encrypted connections to Grafana. To use this feature, specify the following variables in your playbook:

- **metrics_grafana_certificates** to use the **certificate** RHEL system role to generate new certificates on the managed nodes
- **metrics_grafana_cert** and **metrics_grafana_private_key** to specify the path to an existing certificate and private key on the managed nodes
- **metrics_grafana_cert_src** and **metrics_grafana_private_key_src** to copy an existing certificate and private key from the control node to the managed nodes

Jira:RHEL-136607[1]

The **VersionAddendum** option is available in SSH configuration

With this update, you can configure the **VersionAddendum** option in SSH settings for match blocks, host blocks, and global client configurations. This enhancement ensures compatibility with the latest OpenSSH versions and provides granular control over your SSH connections.
Jira:RHEL-138277

The **sshd** system role supports **GSSAPIDelegateCredentials**

The new **GSSAPIDelegateCredentials** parameter provides Generic Security Services Application Programming Interface (GSSAPI) credential delegation in Kerberos environments and enables a seamless single sign-on experience.
As a result, you can automate the configuration of GSSAPI credential delegation to simplify network authentication.

Jira:RHEL-144495

The postgresql RHEL system role now supports PostgreSQL 18

The postgresql RHEL system role, which installs, configures, manages, and starts the PostgreSQL server, now supports PostgreSQL 18.
For more information about this system role, see Installing and configuring PostgreSQL by using the postgresql RHEL system role.

Jira:RHEL-144914[1]

SELinux supports the DCCP and SCTP protocols

With this update, you can manage SELinux port types for Datagram Congestion Control Protocol (DCCP) and Stream Control Transmission Protocol (SCTP). By configuring SELinux port labels for these protocols, you can apply granular access controls and improve system security.
Jira:RHEL-145214

RHEL System Roles support for immutable systems (**ostree**)

You can use RHEL system roles to build and manage immutable operating systems. This provides a consistent management interface across different backend technologies, including **ostree**.
As a result, you can deploy and configure immutable systems using the same roles used for traditional systems, ensuring environment consistency. Note: This feature is currently not compatible with the **nbde_client** role.

Jira:RHELDOCS-21216

In-place upgrade phases automation with the **analysis**, **remediate**, and **upgrade** Ansible roles

With this release, you can use the **analysis**, **remediate**, and **upgrade** Ansible roles to automate the pre-upgrade and upgrade phases of the in-place upgrade. By using these Ansible roles, you can quickly and efficiently upgrade large numbers of systems, saving you time.
For more information, see Upgrading large deployments by using Ansible roles .

Jira:RHEL-141757

6.19. VIRTUALIZATION

Review new features and enhancements for virtualization in Red Hat Enterprise Linux 10.2.

VMs on IBM Z can now use multiple network boot entries

With this update, virtual machines (VMs) on RHEL 10 hosts that use IBM Z hardware can have multiple kernel entries for the **virtio-net** device. As a result, you can use multiple VM kernel boot entries available over PXE if the primary boot device is not bootable.

Jira:RHEL-71834

Faster updates for cryptographic coprocessors on IBM Z

After using the **virsh nodedev-update** command to update a cryptograpic coprocessor (**vfio-ap**) device on an IBM Z host, the new configuration now takes effect significantly faster.

Jira:RHEL-73000[1]

CPI for virtual machines on IBM Z

Virtual machines (VMs) on RHEL 10 hosts that use IBM Z hardware can now use the Control Program Identification (CPI) feature. By using CPI, you can obtain system information about VMs without accessing them. For more information about CPI, see IBM documentation.

Note that on VMs that use IBM Secure Execution, CPI is disabled by default to ensure confidentiality, and must be enabled manually. For instructions, see Setting up IBM Secure Execution on IBM Z.

Jira:RHEL-73008[1]

Configure hostname and FQDN options in **libvirt** XML configuration

The **libvirt** virtualization API supports setting hostname and Fully Qualified Domain Name (FQDN) options for virtual machines on network interfaces that use the **passt** backend. This feature integrates **passt** DHCP and DHCPv6 capabilities to simplify network identity assignment. As a result, you can configure hostname and FQDN directly in the domain XML. For example:

```
<backend type='passt' hostname='vm1' fqdn='vm1.kubevirt.org.'/>
```

Both attributes are optional.

Jira:RHEL-79806

Backup jobs now keep VMs active during guest OS shutdown

Backup jobs initiated through tools such as **virsh backup-begin** now keep the virtual machine (VM) process active even if the guest operating system (OS) shuts down during the operation. Previously, a guest OS shutdown caused **libvirt** to terminate the VM process, which failed the backup and required a manual restart. This enhancement ensures that backup jobs complete successfully regardless of the guest OS state, providing greater reliability and eliminating manual intervention.

Jira:RHEL-80679

The **virtio-win** package introduces the **viosock** driver for Windows virtual machines

Virtual Socket (**vsock**) is a communication interface for direct socket-based communication between a host and virtual machines (VMs) running on the host. With this update, the **virtio-win** package includes the **viosock** driver, which implements **vsock** support in Windows VMs running on a KVM host. The driver enables use cases such as running commands in a Windows VM directly from the host.

The **virtio-win** package also includes the **VsockTcpBridge** service, which provides a **vsock**-to-TCP bridge. This bridge allows existing TCP-based applications in the Windows VM to communicate over the **vsock** interface without modification.

The **viosock** driver is available in the **virtio-win** ISO and installer. When you install the driver, the **VsockTcpBridge** service and the **vsock** provider are configured automatically.

Jira:RHEL-91040

New s390-ccw-virtio-rhel10.2.0 machine type available for IBM Z VMs

The updated **qemu-kvm** package provides a new **s390-ccw-virtio-rhel10.2.0** machine type for IBM Z virtual machines (VMs). This machine type enables Control Program Identification (CPI) and performance-enhanced PCI translation for passthrough PCI devices by default. As a result, IBM Z VMs that use the **s390-ccw-virtio-rhel10.2.0** machine type benefit from improved performance with passthrough PCI devices and CPI without additional configuration.

Jira:RHEL-104009[1]

Block device I/O limits included in libvirt domstats

The **virsh domstats --block** command displays block device I/O limits for virtual machine (VM) block nodes. The limits include:

- Maximum I/O request size
- Maximum I/O vector count
- Memory alignment values

  By using this feature you can inspect the I/O limits that QEMU uses for storage back ends and determine whether your VM configuration is optimal. As a result, you can better debug performance issues and detect incorrect storage configurations.

Jira:RHEL-118671

PCCS for Intel TDX

This update introduces the Provisioning Caching Certification Service (PCCS) for Intel Trust Domain Extensions (TDX). This provides the local caching required to use Intel hosted Provisioning Certification Services (PCS) at scale, and also makes it possible to perform TDX attestation on host systems that are isolated from the public internet.

Jira:RHEL-121612

libvirt introduces a host-model mode for Hyper-V Enlightenments

The **libvirt** package provides a new **host-model** mode for Hyper-V Enlightenments, which automatically enables all Hyper-V enlightenments supported on the host. This mode eliminates the need for separate configuration templates for Intel and AMD hosts. As a result, you can configure **<hyperv mode='host-model'/>** in the XML definition of a virtual machine to automatically apply all host-supported Hyper-V Enlightenments without maintaining separate configurations for each vendor.

Jira:RHEL-122932[1]

Encryption for libvirt secrets

This update introduces the **virt-secrets-init-encryption** service, which encrypts **libvirt** secrets, such as keys for the virtual Trusted Platform Module (vTPM). By default, this encryption uses **systemd** credentials sealing. However, you can use the new **/etc/libvirt/secret.conf** file to specify a custom

key for encrypting secrets, as well as to disable automatic encryption of secrets. As a result, critical vTPM metadata is protected from unauthorized access on the host file system. This also hardens the overall security of the virtualization environment.

Jira:RHEL-7125[1]

Native FUA support for QEMU

With this update, the QEMU emulator no longer needs to emulate the Forced Unit Access (FUA) I/O method, and instead can use FUA natively. This can improve the overall performance of virtual storage, particularly in database workloads.

Jira:RHEL-66064[1]

6.20. SUPPORTABILITY

Review new features and enhancements for supportability in Red Hat Enterprise Linux 10.2.

Ceph mon sessions added to sos report

In the latest version of the **sos** tool, system administrators can effortlessly retrieve a list of active mon sessions from a Ceph cluster. This was accomplished by connecting to the admin socket and executing the **ceph tell mon sessions** command. This feature was implemented to enhance the efficiency of troubleshooting Ceph related problems.
As a result, users can now investigate issues related to Ceph sessions with the data included in a SOS archive.

Jira:RHEL-103783

The new aws plugin in sos collects metadata information

With this update, **sos** includes a plugin that collects metadata information from AWS instances. This update introduces the following notable enhancements:

- Enhances metadata collection from AWS instances by using an **sos** plugin.
- Improves the data gathering process in the **sos** RPM package across RHEL versions.
- Provides an accurate and detailed analysis of AWS instances within the **sos** report.

Jira:RHEL-114887

Improved AAP plugins for more useful diagnostics

Before this update, the **sos** report was collected on **AAP**. With this update, the notable enhancements to the following AAP plugins are:

- **aap_containerized**: Resolved an issue that incorrectly enabled **aap_containerized** on the RPM-based Private Automation Hub servers.
- **aap_controller**: Expanded the set of gathered command outputs and conditionally collect **run_wsbroadcast** or **run_wsrelay** depending on the AWX release version.
- **aap_eda**: Collected service output details based on the installed EDA version. Starting from AAP 2.5, specific commands are used to obtain service status information.

- **aap_gateway**: Added additional command outputs for improved troubleshooting on Gateway servers.
- **aap_hub**: Centralized the collection of service information for PAH servers under a single location within the plugin directory.

Jira:RHEL-121524

SSL certificate control in SOS clean process is available

With this update, you can manage SSL/TLS certificates that contain sensitive data during the SOS clean process. The new **--treat-certificates** option provides the option to remove, obfuscate, or maintain the original binary format of these certificates ensuring that no sensitive data persists. As a result, you can enhance data security and privacy by selecting the treatment for SSL/TLS certificates during the SOS clean process.

Jira:RHEL-142619

Automatic user detection for AAP container runners in SOS reports

With this update, the **sos** utility automatically detects the user running containers for Ansible Application Platform (AAP) deployments. This eliminates the need for manual specification, ensuring the collection of all necessary AAP data.
Jira:RHEL-140738

6.21. CONTAINERS

Review new features and enhancements for containers in Red Hat Enterprise Linux 10.2.

Podman switches to Sequoia-PGP for OpenPGP signatures in RHEL 10

With this update, Podman supports a Sequoia-PGP-based back end for OpenPGP image signatures. Previously, Podman used **GnuPG** (through **gpgme**/**pgpme** bindings) for **OpenPGP** operations. This update includes the following enhancements:

- Verification: The back end is switched from GnuPG to Sequoia-PGP.
- Signing: The current GnuPG workflows continue to exist. With the new **--sign-by-sq-fingerprint** option you can use Sequoia and Sequoia-available keys. The current GnuPG workflows remain supported.
- Algorithm support: Supports modern and post-quantum capable algorithms such as ML-DSA-87+Ed448.
- Improved Skopeo compatibility with FIPS certification.

Jira:RHEL-56365[1]

container-selinux rebased to version 2.244.0-1

The **container-selinux** package, which provides necessary SELinux policies, types, and rules to confine and secure container runtimes, has been rebased to version 2.244.0-1. This version provides important bug fixes and enhancements, most notably:

- Enhanced data protection ensures confidentiality in deployments, while reducing potential

security risks associated with public storage endpoints.

- Errors in package NVR no longer cause reproducible crashes, improving system stability.

Jira:RHEL-111947

gvisor-tap-vsock rebased to 0.8.7-1

The **gvisor-tap-vsock** package, which provides a user space networking stack for virtual machines, particularly those used with Podman, is rebased to upstream version 0.8.7-1. This version provides important fixes and enhancements, most notably, users can integrate a private image registry within a private Microsoft Azure cluster, enhancing security and efficiency of image management.
As a result, the ability to create customizable, secure storage endpoints within the deployment, streamlining storage resource management and reducing potential security risks.

Jira:RHEL-111948

buildah rebased to 1.41.8-1

The **buildah** package, which provides a daemonless command-line tool for building Open Container Initiative (OCI-compliant), is rebased to upstream version 1.41.8-1. This version provides important fixes and enhancements, most notably, you can integrate a private image registry within a private Microsoft Azure cluster, enhancing the management and deployment of container images in a secure and scalable environment.
As a result, a more secure storage solution is available because you can now secure the storage endpoints privately on Azure, protecting their data from unauthorized access. Simplified management of storage endpoints also makes it easier for you to maintain their storage infrastructure.

Jira:RHEL-114411

crun is rebased to 1.25.1-1

The **crun** package provides a fast, lightweight, and low memory Open Container Initiative (OCI) runtime acting as the default, high-performance alternative to **runc** for executing containers. The **crun** is rebased to upstream version 1.25.1-1. This version provides important fixes and enhancements, most notably the following:

- Users can create and manage their own private container registries within a secure Microsoft Azure Kubernetes Service (AKS) cluster. This enhancement streamlines navigation, increases efficiency, and ensures data security.

- Users can deploy and manage their containerized applications with improved security and scalability, enabling seamless integration of third party applications and expanding the functionality of the platform.

- By automating routine tasks, it saves valuable time and effort, allowing them to focus on more complex tasks, improving overall efficiency and productivity.

Jira:RHEL-114419

python-podman rebased to 5.7.0-1

The **python-podman** package is rebased to upstream version 5.7.0-1. With **python-podman**, you can manage Podman containers, images, volumes, and pods. The new version provides important fixes

and enhancements, most notably, you can integrate a private image registry within a secure Azure cluster. The private registry installation ensures a more secure deployment of applications, as it offers enhanced protection for sensitive images.
Jira:RHEL-114423

Unified configuration available for rootless Podman

With this update, rootless Podman introduces a unified system-wide configuration file that enables centralized policy management, a consistent security baseline, and operational standardization across all users.
As a result, you can inherit sensible defaults without manual configuration while maintaining the flexibility to override system defaults through personal configuration files. Additionally, this update ensures backward compatibility, so existing workflows and configurations remain unchanged.

Jira:RHEL-126644

The Container Tools packages have been updated

The updated Container Tools RPM meta-package, which includes the Podman, Buildah, Skopeo, crun, and runc tools, is available. The Buildah package has been updated to version 1.43.1, and Skopeo has been updated to version 1.22.2. Podman release 5.8.2 contains the following notable bug fixes and enhancements over the previous version:

- The **podman machine init --image** command can run **PowerShell-escaped** commands from the user-specified image path in a PowerShell session on the host when you use it on Windows with the Hyper-V backend (CVE-2026-33414).
- Automatic migration from BoltDB to SQLite after a reboot no longer performs a partial migration, leaving some containers in SQLite and others in BoltDB, when Quadlets are in use.
- The **podman quadlet install** command installs files that contain multiple separate Quadlet files. You must separate the files with a **--- delimiter** on a new line and begin each section with a **# FileName=<name>** line to name the new Quadlet.
- The **Quadlet .container** files include the **AppArmor** key to configure a container's AppArmor profile.
- Podman automatically attempts to migrate earlier BoltDB databases to SQLite when the system reboots. This is necessary because the Podman 6.0 release removes support for BoltDB. If automatic migration is not possible, you can manually force a migration with the new **podman system migrate --migrate-db** option.
- Podman loads the path from the VM's filesystem when you run the **podman artifact add** command against a Podman machine VM. This improves performance if you share the path you load or build into the VM instead of streaming the data through the REST API.
- The **podman update** command has a new option, **--ulimit**, to update container ulimits.
- You can use the new **--no-session** option with the **podman exec** command to disable tracking of the exec session, which improves performance and startup time.
- Containers with the **unless-stopped** restart policy restart after a reboot when you enable the **podman-restart.service** service.
- In the **Quadlet.container** file:

    - You can set **Entrypoint=""** to clear the container's entrypoint.
    - A **HealthCmd** supports commands with double-quotes and ensures a functional health check.
    - The **RequiresMountsFor** field correctly handles bind-mount paths that contain spaces.
- Inspecting containers in host network mode no longer causes FreeBSD systems to panic.
- The Libpod System Check endpoint no longer performs operations with bad data after it returns a 400 error.
- The remote attach API for containers (Libpod & Compat) no longer panics due to a rare race condition.
- The system no longer improperly adds options from the default driver, which previously prevented the Secret Create API from creating functional secrets using the shell driver. You can enter the secret directly at the terminal with the **podman secret create** command instead of providing it through a pipe.
- Added new APIs for interacting with Quadlets:
    - **GET**/libpod/quadlets/{name}/file`: Print the contents of a Quadlet file.
    - **GET**/libpod/quadlets/{name}/exists`: Check if the given Quadlet exists.
    - **POST /libpod/quadlets**: Install one or more Quadlets.
    - **DELETE /libpod/quadlets**: Remove one or more Quadlets.
    - **DELETE /libpod/quadlets/{name}**: Remove a single Quadlet.
- Containers created by the **podman play kube** command no longer run health checks before the **initialDelaySeconds** option expires, and the **podman kube play** command correctly handles precedence between environment variables set by both the **envFrom** and **env** fields.
- The **podman build** command's **--pull=newer** option functions correctly.
- The **podman artifact push** and **podman artifact pull** commands no longer ignore authentication credentials given by the **--authfile** option.
- The **podman run --pod-id-file** option is properly validated, preventing the creation of containers in pods with improper user namespace configuration.
  For more information about notable changes, see Upstream release notes.

Jira:RHEL-127903

The **fuse-overlayfs** rebased to 1.16-1

The **fuse-overlayfs** package, a user space implementation of the OverlayFS file system provides rootless containers, which Podman or Buildah run, is rebased to upstream version 1.16-1. This version provides important fixes and enhancements, most notably the following:

- Updated database connection settings resolve intermittent connection errors, making error logging functional, and ensuring smooth operation and reduced downtime for users.

Jira:RHEL-128521

Support for updates in air-gapped and disconnected environments

This update introduces air-gapped and disconnected updates for RHEL deployments, enabling edge deployments to perform updates without internet connectivity. As a result, you can benefit from greater flexibility and reliability for offline updates, improving deployment management in remote or secure environments.

Jira:RHELDOCS-20708[1]

Signing container images by using Sequoia-PGP is available

With this update, Podman supports a Sequoia-PGP-based backend for OpenPGP image signatures. Previously, Podman used **GnuPG** (**gpgme**/**pgpme** bindings) for **OpenPGP** operations. This update includes the following enhancements:

- Verification: the backend is switched from GnuPG to Sequoia-PGP.
- Signing: the current GnuPG workflows continue to exist. New **--sign-by-sq-fingerprint** option allow you to use Sequoia and Sequoia-available keys. Current GnuPG workflows remain supported.
- Algorithm support: Supports modern and post-quantum capable algorithms such as ML-DSA-87+Ed448.

Jira:RHELDOCS-21869[1]

New container images are available

The **rhel10**/**ruby-40**, **rhel10**/**postgresql-18**, **rhel10**/**python-314-minimal**, **rhel10**/**mariadb-118** and **rhel10**/**php-84** container images are now available in the Red Hat Container Registry. The notable enhancements for each image are:

- **rhel10**/**ruby-40**: You use the Ruby 4.0 container as your base platform to build and run diverse Ruby 4.0 applications and frameworks. This container image includes the npm utility, so you can install JavaScript modules for your web applications.
- **rhel10**/**postgresql-18**: You can use this container image to package the PostgreSQL **postgres** daemon and client application in a container. The **postgres** server daemon accepts your connections from clients and provides you access to content from PostgreSQL databases.
- **rhel10**/**python-314-minimal**: You use the full container image as a universal base image to build your containerized applications. However, this universal nature means that the resulting containers consume a lot of disk space. This happens mainly because the image contains npm, compilers, header files, and other packages you might need to install and deploy your applications.
- **rhel10**/**mariadb-118**: You use this container image to package the MariaDB **mysqld** daemon and client application into a container. The **mysqld** server daemon accepts your client connections and provides you with access to content from MySQL databases.
- **rhel10**/**php-84**: You can use this container image as a base platform for building and running various PHP 8.4 applications and frameworks. You can also install JavaScript modules for the web applications. This container image includes an npm utility.

Jira:RHELDOCS-21963

6.22. RHEL LIGHTSPEED

Review new features and enhancements for RHEL Lightspeed in Red Hat Enterprise Linux 10.2.

Color support for the command-line assistant

With this update, the command-line assistant supports color output by default, aligning its appearance with other RHEL command-line tools. This update improves output readability through increased visual contrast.

You can disable color output by using the **--plain** option or by setting the **NO_COLOR=1** environment variable.

Jira:RHELDOCS-21814[1]

SAP Solutions documentation added to RHEL Lightspeed

With this enhancement, RHEL Lightspeed includes the Red Hat Enterprise Linux for SAP Solutions documentation set in its knowledge base. You can now ask RHEL Lightspeed technical questions specific to SAP deployments on RHEL. This update provides more accurate and context-aware responses for SAP-related administrative and configuration tasks.

Jira:RHELDOCS-21815[1]

CHAPTER 7. TECHNOLOGY PREVIEW FEATURES

Review newly identified and previously known Technology Preview features available in Red Hat Enterprise Linux 10.2.

For information about Red Hat scope of support for Technology Preview features, see *Technology Preview Features Support Scope*.

Additional resources

- Technology Preview Features Support Scope

7.1. INSTALLER AND IMAGE CREATION

Review new Technology Preview features available for installer and image creation in Red Hat Enterprise Linux 10.2.

A new **bootc** Kickstart command available as a Technology Preview

The installer includes support for the new **bootc** Kickstart command as a Technology Preview. It enables the deployment of bootable containers. Although comparable to the existing **ostreecontainer** Kickstart command, this implementation relies on the **bootc** utility to manage both operating system content provisioning and boot loader setup. You can use the command in the following way:

```
# bootc --source-imgref=<transport>:<registry>/<namespace>/<name>:<tag> --target-imgref=<registry>/<namespace>/<name>:<tag>
```

NOTE

The feature does not support complex partitioning schemes that span multiple storage devices or custom mount point configurations beyond standard layouts.

For more information, see Installing RHEL from a bootable container image by using the bootc Kickstart command.

Jira:RHEL-58215[1]

Integrity Image Sealing is available as a Technology Preview

With this Technology Preview, you can cryptographically seal your boot container images by using your organization's Secure Boot keys, ensuring complete operating system integrity from build through runtime. This is based on Unified Kernel Images that embed a digest of the target container root filesystem, alongside a bootloader (such as **systemd-boot**) also signed with your key.

As a result, you can achieve higher security than current solutions and meet compliance requirements for tamper-proof systems, ensuring the integrity of the executed code from hardware to the operating system. The container image includes a Unified Kernel Image and covers the integrity of the boot process, path, and host operating system. For more information, see Building sealed images.

Jira:RHELDOCS-20426[1]

7.2. FILE SYSTEMS AND STORAGE

Review new Technology Preview features available for file systems and storage in Red Hat Enterprise Linux 10.2.

TLS support for NVMe-TCP as a Technology Preview in NVMe/TLS

NVMe/TLS, available as a Technology Preview, complies with standard TLS key derivation specifications. This update introduces a breaking change to TLS Pre-Shared Key (PSK) import functionality. This change affects the **gen-tls-key** and **check-tls-key** commands for **nvme-cli** versions earlier than 2.16 and **libnvme** versions earlier than 1.16.

If NVMe/TLS connections to a storage target fail after an upgrade, perform one of the following actions:

- Use the **--compat** flag with **nvme-cli** when you import TLS PSKs to maintain operations with existing out-of-spec implementations.
- If connections still fail when you use the **--compat** flag after a storage target upgrade, you must re-provision the TLS PSKs to match the vendor's updated implementation.

Jira:RHEL-135994

7.3. IDENTITY MANAGEMENT

Review new Technology Preview features available for Identity Management (IdM) in Red Hat Enterprise Linux 10.2.

Passwordless authentication mechanisms are available in GDM (Technology Preview)

Identity Management (IdM) administrators can configure the GNOME Display Manager (GDM) login screen to display multiple authentication mechanisms. In addition to existing smart card authentication, administrators can enable new passwordless methods, such as external identity providers (EIdP) and FIDO2-compatible passkeys. Enable the **with-switchable-auth** feature in **authselect** and configure the System Security Services Daemon (SSSD) to allow users to choose their preferred credential directly at login.

Passwordless authentication aligns with zero trust architecture by replacing static passwords with cryptographic proof that verifies both user identity and device integrity for each access request. For detailed configuration instructions and a list of current limitations, see Enabling authentication mechanism selection in GDM using SSSD.

Jira:RHEL-11913[1]

The IdM Modern Web UI is available (Technology Preview)

With this update, Identity Management (IdM) provides the Modern Web UI as a Technology Preview. This new interface features updated design and is available at the **/ipa/modern-ui** endpoint. You can access the new interface through a link on the IdM Web UI login screen.

As a Technology Preview, the Modern Web UI is under active development and intended for experimentation in non-production environments. Provide feedback at the FreeIPA Web UI community project to help improve the interface.

Jira:RHEL-90121

7.4. SSSD

Review new Technology Preview features available for SSSD in Red Hat Enterprise Linux 10.2.

SSSD supports generic Identity Provider integration (Technology Preview)

SSSD provides a generic identity provider (IdP), initially supporting Keycloak and Entra ID. You can configure SSSD to read users and groups directly from these IdPs and authenticate users by using the OAuth 2.0 Device Authorization Grant (RFC 8628). This allows you to use modern IdPs for centralized authentication and access management. This capability is a Technology Preview feature. For more information, see the **sssd-idp(5)** man page.

Jira:RHEL-4990

7.5. DESKTOP

Review new Technology Preview features available for desktop in Red Hat Enterprise Linux 10.2.

Interactive authentication selection is available on the GDM Login Screen (Technology Preview)

The GNOME Display Manager (GDM) provides an interface for users to select a preferred authentication method. Previously, the graphical login environment restricted users to a single authentication method. With this update, users can switch between methods such as external identity providers (EIdP), FIDO2-compatible passkey devices, or smart cards directly from the login screen. The feature is available as a Technology preview.
For more information to enable this functionality and a list of current limitations, see Enabling authentication mechanism selection in GDM using SSSD.

Jira:RHEL-14524[1]

mutter rebase introduces an HDR switch for HDR displays (Technology Preview)

The **mutter** 49 rebase introduces a High Dynamic Range (HDR) switch in the display settings. The HDR switch enables users to change between HDR and Standard Dynamic Range (SDR) modes, which improves media and graphics visuals on compatible devices. This feature is available as a Technology Preview.
Jira:RHEL-144935

7.6. VIRTUALIZATION

Review new Technology Preview features available for virtualization in Red Hat Enterprise Linux 10.2.

Secure Boot for VMs on ARM64 (Technology Preview)

As a Technology Preview, you can now configure the Secure Boot feature for virtual machines (VMs) on RHEL 10 hosts that use ARM64 hardware (also known as AArch64). Secure Boot ensures that the VM is running a cryptographically signed operating system (OS). This can be useful if the guest OS of a VM has been altered by malware. In such a scenario, Secure Boot prevents the VM from booting, which stops the potential spread of the malware to your host machine.
Jira:RHEL-82645

Live migration for S3-PR (Technology Preview)

As a Technology Preview, you can now live migrate a virtual machine (VM) with enabled SCSI3-Persistent Reservation (S3-PR), with the reservation state being preserved after the migration. To do this, you must use the following XML configuration for the VM:

```
<reservations managed="no" migration="yes">
```

Note, however, that migrating a VM with S3-PR and this configuration to a host that uses a previous version of QEMU fails.

Jira:RHEL-135115

SEV-SNP is available on RHEL hosts as a Technology Preview

As a Technology Preview, you can enable Secure Encrypted Virtualization-Secure Nested Paging (SEV-SNP) on RHEL hosts. SEV-SNP is a hardware-based security feature that provides strong memory encryption and integrity protection for virtual machines, isolating them from the hypervisor and other system software.
SEV-SNP is available only with AMD CPUs, and you must use the **snphost** package to configure the feature on the host.

Jira:RHELDOCS-19757[1]

7.7. CONTAINERS

Review new Technology Preview features available for containers in Red Hat Enterprise Linux 10.2.

krun runtime for RHEL Container Workloads is a Technology Preview

Red Hat Enterprise Linux offers the **krun** runtime as a Technology Preview for running container workloads. You can launch containers inside lightweight microVMs, which provides an additional isolation boundary for your workloads by using the **crun** configured to support **krun**. This feature improves container workload performance, security, and addresses an issue where running containers by using **krun** fails because RHEL did not previously include a version of the **krun** runtime.
Jira:RHEL-161090

7.8. TECHNOLOGY PREVIEW FEATURES IDENTIFIED IN RHEL 10.1

Review Technology Preview features that were introduced in Red Hat Enterprise Linux 10.1.

For information about Red Hat scope of support for Technology Preview features, see *Technology Preview Features Support Scope*.

Additional resources

- Technology Preview Features Support Scope

7.8.1. Installer and image creation

Review Technology Preview features introduced for installer and image creation in Red Hat Enterprise Linux 10.1.

image-builder-cli replaces **osbuild-composer** and **composer-cli** (Technology Preview)

With this release, you can install and use the new **image-builder-cli** package to build an image with one command. The new tool supports containers and enhances your user experience to create a container image that you can use to build other images. This capability is a Technology Preview feature. For more details, see Installing RHEL image builder .

Jira:RHELDOCS-20354[1]

7.8.2. Shells and command-line tools

Review Technology Preview features introduced for shells and command-line tools in Red Hat Enterprise Linux 10.1.

RHEL 10.1 provides ReaR on aarch64 (Technology Preview)

RHEL 10.1 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-84286[1]

7.8.3. Dynamic programming languages, web and database servers

Review Technology Preview features introduced for dynamic programming languages, web and database servers in Red Hat Enterprise Linux 10.1.

Node.js 24 is available as a Technology Preview

A new **nodejs24** component is available as a Technology Preview in Red Hat Enterprise Linux 10.1. This update introduces Node.js 24, which includes new features, bug fixes, security updates, and performance improvements compared to Node.js 22 in RHEL 10.0.

Currently, the **nodejs24** package provides versioned binaries (**/usr/bin/node-24**, **/usr/bin/npm-24**, and **/usr/bin/npx-24**). To use these binaries, update the shebang lines in your scripts to reference the version-specific paths. The ability for **nodejs24** to provide the base binaries (**/usr/bin/node** and related files) might be included in a future update.

To install the **nodejs24** package, enter:

```
# dnf install nodejs24
```

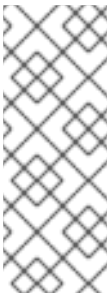

NOTE

On Red Hat Enterprise Linux 10, configure FIPS mode during installation. Switching the methods after installing RHEL is documented only for RHEL 9 does not apply to RHEL 10. In the RHEL build of Node.js 24, downstream patches prevent the use of the **--force-fips** runtime flag. Passing **--force-fips** results in an error regardless of the system FIPS state. If you encounter an error that links to RHEL 9 documentation, note that those steps do not work on RHEL 10.

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

Jira:RHEL-90826

7.8.4. Identity Management

Review Technology Preview features introduced for Identity Management (IdM) in Red Hat Enterprise Linux 10.1.

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]

7.8.5. Virtualization

Review Technology Preview features introduced for virtualization in Red Hat Enterprise Linux 10.1.

Virtual Socket to TCP bridge is available as a Technology Preview

As a Technology Preview, you can use a Virtual Socket (**vsock**) to TCP bridge. By using this bridge, you can securely expose a virtual machine (VM) service, such as SSH, to the host machine without configuring any IP networking.
To bridge your host's connection directly to the SSH service inside the VM over the hypervisor's private **vsock** channel, you can use a relay tool such as **socat**.

Jira:RHEL-91041

CCA on ARM virtual machines is available as a Technology Preview

As a Technology Preview, you can enable Confidential Compute Architecture (CCA) on RHEL 10.1 and later virtual machines (VMs). CCA, built on top of Realm Management Extension (RME), helps to maintain data privacy while it is in use within a virtual machine.
Currently, CCA can only be enabled on ARM VMs as a Technology Preview and not on a RHEL host.

Jira:RHEL-83042

TDX is available on RHEL hosts as a Technology Preview

As a Technology Preview, you can enable Trust Domain Extensions (TDX) on RHEL hosts. TDX is a hardware-based security feature that provides strong memory encryption and integrity protection for virtual machines, isolating them from the hypervisor and other system software.
TDX is available only with Intel CPUs.

Jira:RHEL-111863[1]

7.8.6. Containers

Review Technology Preview features introduced for containers in Red Hat Enterprise Linux 10.1.

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-88122

7.9. TECHNOLOGY PREVIEW FEATURES IDENTIFIED IN RHEL 10.0

Review Technology Preview features that were introduced in Red Hat Enterprise Linux 10.0.

For information about Red Hat scope of support for Technology Preview features, see *Technology Preview Features Support Scope*.

Additional resources

- Technology Preview Features Support Scope

7.9.1. Software management

Review Technology Preview features introduced for software management in Red Hat Enterprise Linux 10.0.

Support for signing packages with Sequoia PGP (Technology Preview)

The **macros.rpmsign-sequoia** macro file that configures RPM to use Sequoia PGP instead of GnuPG for signing packages is now available as a Technology Preview. To enable its usage, perform the following steps:

1. Install the following packages:

```
# dnf install rpm-sign sequoia-sq
```

2. Copy the **macros.rpmsign-sequoia** file to the **/etc/rpm/** directory:

```
$ cp /usr/share/doc/rpm/macros.rpmsign-sequoia /etc/rpm/
```

Jira:RHEL-56363[1]

7.9.2. Networking

Review Technology Preview features introduced for networking in Red Hat Enterprise Linux 10.0.

WireGuard VPN (Technology Preview)

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

Jira:RHELDOCS-20056[1]

KTLS (Technology Preview)

In RHEL, Kernel Transport Layer Security (KTLS) is provided as a Technology Preview. KTLS handles TLS records by 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.

Jira:RHELDOCS-20440[1]

7.9.3. Kernel

Review Technology Preview features introduced for kernel in Red Hat Enterprise Linux 10.0.

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]

7.9.4. File systems and storage

Review Technology Preview features introduced for file systems and storage in Red Hat Enterprise Linux 10.0.

ublk_drv driver is available as a Technology Preview

The **ublk_drv** kernel module is now enabled as a Technology Preview. It provides the **ublk** framework with which you can create and build high-performance block devices from userspace. Currently, **ublk** requires userspace implementations, such as the Userspace Block Driver (**ublksrv**) or the Rust-based **ublk** (**rublk**), to function effectively.

Jira:RHELDOCS-19891[1]

NVMe/TCP using TLS is available (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 10.0. For instructions, see Configuring an NVMe/TCP host using TLS with Pre-Shared-Keys .

Jira:RHELDOCS-19968[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-20041[1]

Limited shrinking of XFS file systems is available as Technology Preview

You can reduce the size of XFS file systems by using the **xfs_growfs** utility as a Technology Preview. You can remove blocks from the end of the file system by using **xfs_growfs**, provided that all of the following conditions are true:

- No metadata or data is allocated within the range to be removed.
- The requested size is within the last allocation group.

Jira:RHELDOCS-20042[1]

Mounting XFS file systems with blocks larger than system page is available as Technology Preview

You can now mount XFS file systems created with a block size larger than the system page size as a Technology Preview. For example, a file system with 16-KB blocks can now be mounted on a system with a 4-KB page size, such as x86_64.

Jira:RHELDOCS-20043[1]

io-uring interface is available as a Technology Preview

The **io_uring**, which is an asynchronous I/O interface, is available as a Technology Preview. By default, this feature is disabled in RHEL 10. You can enable this interface by setting the **kernel/io_uring_disabled** variable:

- For all users:

```
# echo 0 > /proc/sys/kernel/io_uring_disabled
```

- For root only:

```
# echo 1 > /proc/sys/kernel/io_uring_disabled
```

You can also disable **io_uring** for all processes:

```
# echo 2 > /proc/sys/kernel/io_uring_disabled
```

Jira:RHEL-65347

NVMe/TCP Boot with NBFT is available as a Technology Preview

NVMe/TCP Boot by using the NVM Express Boot Specification (NBFT) is available on select server platforms as a Technology Preview. Consult your server manufacturer for platform-specific details and compatibility information.

Jira:RHELDOCS-21587[1]

7.9.5. Compilers and development tools

Review Technology Preview features introduced for compilers and development tools in Red Hat Enterprise Linux 10.0.

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]

7.9.6. Identity Management

Review Technology Preview features introduced for Identity Management (IdM) in Red Hat Enterprise Linux 10.0.

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 on IdM servers and replicas, and the **ipa-client-encrypted-dns** package on IdM clients. Administrators can enable DoT during the installation by 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.

The following 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 the **relaxed** DNS policy, which allows fallback to unencrypted DNS. You can enforce encrypted-only communication by 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 by using **ipa-dns-install** with the new DoT options.

See Securing DNS with DoT in IdM for more details.

Jira:RHEL-67912

7.9.7. Virtualization

Review Technology Preview features introduced for virtualization in Red Hat Enterprise Linux 10.0.

AMD SEV, SEV-ES, and SEV-SNP for KVM virtual machines are available as a Technology Preview

As a Technology Preview, RHEL 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 VM security.
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 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 to 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.
- SEV-SNP works only on 3rd generation AMD EPYC CPUs (codenamed Milan) or later.

Also note that RHEL includes SEV, SEV-ES, and SEV-SNP encryption, but not the SEV, SEV-ES, and SEV-SNP security attestation and live migration.

Jira:RHELDOCS-16800[1]

Creating nested virtual machines (Technology Preview)

Nested KVM virtualization is provided as a Technology Preview for KVM virtual machines (VMs) running on Intel, AMD64, and IBM Z hosts with RHEL 10. With this feature, a RHEL 7, RHEL 8, or RHEL 9 VM that runs on a physical RHEL 10 host can act as a hypervisor, and host its own VMs.

Jira:RHELDOCS-20080[1]

New package: **trustee-guest-components** (Technology Preview)

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-73770[1]

7.9.8. Containers

Review Technology Preview features introduced for containers in Red Hat Enterprise Linux 10.0.

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-32266

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 informal, reference the man page.
Jira:RHEL-70218

The vrf option for thepodman network create is available as a Technology Preview

The **podman network create** command now provides the **vrf** value for the **--opt** option, as a Technology Preview. The **vrf** value assigns a virtual routing and forwarding instance (VRF) to the bridge interface. It accepts the name of the VRF and defaults to none.

WARNING

This option can only be used with the Netavark network backend.

Jira:RHEL-89373

7.10. TECHNOLOGY PREVIEW FEATURES IDENTIFIED IN PREVIOUS RELEASES

Review Technology Preview features that were introduced in earlier Red Hat Enterprise Linux versions.

For information about Red Hat scope of support for Technology Preview features, see *Technology Preview Features Support Scope*.

Additional resources

- Technology Preview Features Support Scope

7.10.1. Networking

Review Technology Preview features introduced for networking in previous Red Hat Enterprise Linux versions.

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

7.10.2. Identity Management

Review Technology Preview features introduced for Identity Management (IdM) in previous Red Hat Enterprise Linux versions.

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:RHELDOCS-20690[1]

7.10.3. Virtualization

Review Technology Preview features introduced for virtualization in previous Red Hat Enterprise Linux versions.

VDUSE for RHEL networking is available as a Technology Preview

The virtio Data Path Acceleration (vDPA) device in userspace (VDUSE) feature is now available as a Technology Preview for RHEL networking. VDUSE is a Linux kernel mechanism, which allocates user-space for vDPA devices specifically. This mechanism enables a user-space process to register a **virtio-class** device, such as a NIC or block device, with the kernel in a controlled manner. As a result, you can use it on virtual machines or the host through standard vDPA or virtio interfaces.

Jira:RHEL-76477[1]

CHAPTER 8. DEVELOPER PREVIEW FEATURES

Review Developer Preview features that are available in Red Hat Enterprise Linux 10.2.

For information about Red Hat scope of support for Developer Preview features, see *Developer Preview - Scope of Support*.

Additional resources

- Developer Preview - Scope of Support

8.1. INSTALLER AND IMAGE CREATION

Review Developer Preview features available for installer and image creation in Red Hat Enterprise Linux 10.2.

Unified storage for bootc is available as a Developer Preview

With this Developer Preview feature, you can configure **bootc** to pull the host image into bootc-owned container storage, enhancing container image management efficiency. This feature reduces the need for repeated image pulls on the same host, allowing the host image to be reused by container runtimes, such as podman for running and building layered images. It improves overall network performance by using **zstd** chunking when available and streamlines development and testing processes.
As a result, by configuring the unified storage feature on your system, you can store, manage, and utilize container images more efficiently, providing a seamless bootc container management experience.

Jira:RHELDOCS-21395[1]

8.2. RHEL LIGHTSPEED

Review Developer Preview features available for RHEL Lightspeed in Red Hat Enterprise Linux 10.2.

goose as an alternative for the existing RHEL command-line assistant back-end service

This update introduces **goose** as an alternative to the command-line interface (CLI) client for the existing RHEL command-line assistant (CLA) back-end service. These changes do not affect the **command-line assistant** back end, and functionalities remain the same. The **goose** package is available on the RHEL Extensions repository. You can use the **goose** text chat interface interactively during the chat session.
When you use **goose**, before you execute the MCP server or a local Linux command, it prompts you to confirm that the LLM does not break your system or workflow.

For more information, see the Red Hat Knowledgebase article link:https://access.redhat.com/articles/7142302Optimize the RHEL command-line assistant tasks by using goose-redhat].

IMPORTANT

The existing RHEL command-line assistant (CLA) that uses a Red Hat-provided LLM is still available and integrated directly into the RHEL 10 (and late RHEL 9) ecosystem. It is usually installed by using standard DNF repositories and is a core part of the RHEL Lightspeed offering.

Jira:RSPEED-2846[1]

Agent skills for Red Hat Enterprise Linux are available (Developer Preview)

With this Developer Preview, Red Hat Enterprise Linux introduces two new agent skills designed to enhance AI-driven administration and troubleshooting for Red Hat Enterprise Linux (RHEL). The Agent skills for RHEL are built on the Agent Skills (**SKILL.md**) open standard. These integrations equip AI agents, such as Cursor and Claude Code, with workflows and domain expertise directly from Red Hat.
By using these skills, you can help ensure that the AI-generated guidance aligns with RHEL standards, moving away from generic Linux advice toward recommended Red Hat methodologies.

- Best Practices Skill for RHEL focuses on proactive maintenance, system health, and troubleshooting. It provides the AI tool with a structured framework to assist with diagnosing complex RHEL environments.
- Translator Skill for RHEL is designed for users migrating from other distributions or earlier systems. This skill translates general Linux concepts into RHEL-native equivalents.
  For more information, see Best practices agent skill for Red Hat Enterprise Linux and Translator agent skill for Red Hat Enterprise Linux .

Jira:RHELDOCS-22164[1]

8.3. DEVELOPER PREVIEW FEATURES IDENTIFIED IN RHEL 10.1

Review Developer Preview features that were introduced in Red Hat Enterprise Linux 10.1.

Additional resources

- Developer Preview - Scope of Support

8.3.1. RHEL Lightspeed

Review Developer Preview features introduced for RHEL Lightspeed in Red Hat Enterprise Linux 10.1.

The linux-mcp-server for Red Hat Enterprise Linux is available (Developer Preview)

This Developer Preview introduces the **linux-mcp-server** for Red Hat Enterprise Linux (RHEL), which is designed to bridge the gap between RHEL systems and large language models (LLMs). By using this Model Context Protocol (MCP) server, you can enable AI applications to perform context-aware troubleshooting on RHEL systems, including log and performance analysis. For more details, see Using the MCP server for RHEL to enable AI assistants to run, discover, and troubleshoot complex issues.

Jira:RHELDOCS-21153[1]

CHAPTER 9. REMOVED FEATURES

Review features that were removed in Red Hat Enterprise Linux 10.2.

All removed features were deprecated in earlier releases and are no longer supported. For information regarding functionality that is present in RHEL 9 but has been removed in RHEL 10, see *Considerations in adopting RHEL 10*.

Additional resources

- Considerations in adopting RHEL 10

9.1. SECURITY

Review removed features for security in Red Hat Enterprise Linux 10.2.

Non-post-quantum KEX removed from the FUTURE policy

The **FUTURE** system-wide cryptographic policy no longer allows traditional, non-post-quantum, key exchange (KEX) methods. With this update, you can use only hybrid Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM) key exchange algorithms.
This aligns with the target use of the **FUTURE** policy in controlled, isolated deployments and impacts interoperability with the public internet. For example, when you switch RHEL 10 to the **FUTURE** policy, you cannot connect to the cdn.redhat.com network. Also, Java clients cannot connect to TLS servers running in **FUTURE**.

Note that upcoming RHEL releases might further restrict signature algorithms in certificates and signatures to allow only post-quantum cryptography (PQC) in the **FUTURE** policy.

Jira:RHEL-93769

9.2. INFRASTRUCTURE SERVICES

Review removed features for infrastructure services in Red Hat Enterprise Linux 10.2.

Vim launch from vi command removed

The **vi** command no longer launches the Vim editor when both **vim-minimal** and **vim-enhanced** are installed. Instead, the **vi** command starts the minimal editor from **vim-minimal**. To use Vim, run the **vim** command.

Jira:RHEL-145868[1]

9.3. IDENTITY MANAGEMENT

Review removed features for Identity Management (IdM) in Red Hat Enterprise Linux 10.2.

The SSSD ipa_enable_dns_sites option is removed

The **ipa_enable_dns_sites** option in SSSD is removed and is no longer available in RHEL 10.2. Because the corresponding server-side functionality was never implemented, the option was non-functional. It has been removed to simplify configuration and avoid confusion.
Jira:RHEL-127792

CHAPTER 10. DEPRECATED FEATURES

Review functionalities that are deprecated in Red Hat Enterprise Linux 10.2.

Deprecated functionalities are fully supported, which means that they are tested and maintained, and their support status remains unchanged within Red Hat Enterprise Linux 10. However, they will likely not be supported in a future major version release, and are not recommended for new deployments on the current or future major versions of Red Hat Enterprise Linux.

Features can be deprecated during a major version's release cycle.

Deprecated features are listed in all future release notes until they are removed. For a complete list of deprecated features, see the release notes for the latest minor version. For information about the length of support, see *Red Hat Enterprise Linux Life Cycle* and *Red Hat Enterprise Linux Application Streams Life Cycle*.

Additional resources

- Red Hat Enterprise Linux Life Cycle
- Red Hat Enterprise Linux Application Streams Life Cycle

10.1. HIGH AVAILABILITY AND CLUSTERS

Review deprecated functionalities for high availability and clusters in Red Hat Enterprise Linux 10.2.

SCTP transport for knet is now deprecated in Corosync

Previously, the **knet** transport protocol in Corosync allowed the selection of Stream Control Transmission Protocol (SCTP), although this specific transport was not officially supported in RHEL. With this update, using SCTP for **knet** transport is officially deprecated. The option to use SCTP might be removed in a future release.

As a result, users are advised to transition to supported transport protocols. The **pcs cluster setup**, **pcs cluster link add**, and **pcs cluster link update** commands now display a warning if SCTP is specified for **knet** transport.

Jira:RHEL-126839

10.2. CONTAINERS

Review deprecated functionalities for containers in Red Hat Enterprise Linux 10.2.

MySQL80 and Python 3.11 container images are deprecated

The **MySQL80** and **Python 3.11** container images are now deprecated and will no longer receive feature updates. To maintain support and receive new features, you must migrate to the **MySQL84** and **Python 3.12** container images.

Jira:RHELDOCS-22088[1]

The bootc-image-builder tool is deprecated

The **bootc-image-builder** tool, to convert **bootc** images into disk images for different platforms and formats, is deprecated and might be removed in a future major release. It remains supported for the

lifetime of Red Hat Enterprise Linux (RHEL) 10. You can create bootable containers and disk images by using the RHEL image builder instead.

Jira:RHELDOCS-22154[1]

10.3. DEPRECATED FEATURES IDENTIFIED IN RHEL 10.1

Review functionalities that are deprecated in Red Hat Enterprise Linux 10.1.

10.3.1. Security

Review functionalities deprecated for security in Red Hat Enterprise Linux 10.1.

oqsprovider and liboqs are deprecated

The **oqsprovider** and **liboqs** packages, which provided post-quantum cryptography (PQC) for OpenSSL 3.0, are deprecated and might be removed in a future major release. Instead, use the PQC functionality provided by OpenSSL 3.5.

Jira:RHEL-97489[1]

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-99813

10.3.2. Compilers and development tools

Review functionalities deprecated for compilers and development tools in Red Hat Enterprise Linux 10.1.

GCC Toolset 15 environment script replaces Software Collections (scl enable)

Previously, the **scl enable gcc-toolset-15 <command>** command was used to manage the development environment for GCC Toolset 15 on Red Hat Enterprise Linux. In RHEL 10, Software Collections are no longer used for this purpose. As a consequence, the **scl enable** option does not work with **gcc-toolset-15**.

Use the new **gcc-toolset-15-env** script, which runs the specified command with the GCC toolset environment:

```
gcc-toolset-15-env <command>
```

If a command is not specified, the script opens a default shell (**sh**) in the GCC toolset environment.

As a result, users must use **gcc-toolset-15-env** instead of **scl enable** to access GCC Toolset 15 in RHEL 10.

Jira:RHEL-88743[1]

10.3.3. Virtualization

Review functionalities deprecated for virtualization in Red Hat Enterprise Linux 10.1.

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-89426[1]

The rtl8139 NIC has been deprecated for VMs

With this update, the **rtl8139** network interface controller type has been deprecated, and will become unsupported for use in virtual machines in a future major release of RHEL. If you require using a non-virtio NIC type on your host, use the **e1000** or **e1000e** NIC instead.

Jira:RHEL-45624

10.4. DEPRECATED FEATURES IDENTIFIED IN RHEL 10.0

Review functionalities that are deprecated in Red Hat Enterprise Linux 10.0.

10.4.1. Installer and image creation

Review functionalities deprecated for installer and image creation in Red Hat Enterprise Linux 10.0.

The squashfs package has been deprecated

The **squashfs** package has been deprecated, and will be removed in a future major RHEL release. As an alternative, **dracut** has support for mounting **erofs**.

Jira:RHELDOCS-18903[1]

gdisk is removed from boot.iso in RHEL 10

The **gdisk** partitioning utility is removed from the **boot.iso** image type in RHEL 10. You still can use **gdisk** in your Kickstarts. For the **boot.iso** image type, other tools are available for handling GPT disks, for example, the **parted** utility.

Jira:RHELDOCS-18904[1]

The module Kickstart command has been deprecated

Anaconda has deprecated its support for DNF modularity, and as a consequence the **module** Kickstart command has been deprecated. This might impact you if you are using modules in the **%packages** section of your Kickstart files or the **module** Kickstart command. This change is implemented for simplifying the installation process and ensuring a more consistent experience moving forward.

Jira:RHEL-34829

The inst.gpt boot option is now deprecated

The **inst.gpt** boot option is now deprecated and will be removed in the future releases. To specify a preferred disk label type, use the **inst.disklabel** boot option. Specify **gpt** or **mbr** to create GPT or MBR disk labels.

Jira:RHELDOCS-18491[1]

10.4.2. Security

Review functionalities deprecated for security in Red Hat Enterprise Linux 10.0.

ENGINE API in OpenSSL is deprecated

In RHEL 10, ENGINE API is deprecated and is planned to be removed in a future major release. No new applications should be built by using the ENGINE API. To keep application binary interface (ABI) and existing applications working, OpenSSL still exports the ENGINE symbols. To prevent new applications from using ENGINE API, OpenSSL sets the **OPENSSL_NO_ENGINE** flag system-wide, and the header **engine.h** that exposes the ENGINE API has been removed.
Jira:RHEL-45704

crypto-policies now set **allow-rsa-pkcs1-encrypt = false** for GnuTLS

In RHEL 10, the GnuTLS library blocks encryption and decryption with the RSA PKCS #1 v1.5 padding by default. Except for the LEGACY policy, the **allow-rsa-pkcs1-encrypt = false** option is specified in all system-wide cryptographic policies (DEFAULT, FUTURE, and FIPS).
Jira:RHEL-64746

HMAC-SHA-1 in FIPS mode is deprecated

The HMAC-SHA-1 cryptographic algorithm is deprecated in FIPS mode, and it might be removed in a future release. Outside FIPS mode, support for HMAC-SHA-1 is preserved.
Jira:RHELDOCS-18674

10.4.3. Software management

Review functionalities deprecated for software management in Red Hat Enterprise Linux 10.0.

Modularity is deprecated

In RHEL 10, the modularity functionality is deprecated and will be removed in a future major release. Therefore, the DNF **module** command displays a deprecation warning.

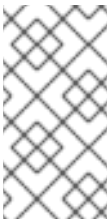

NOTE

In previous RHEL major versions, some Application Streams were available as modules as an extension to the RPM format. In RHEL 10, Red Hat does not intend to provide any Application Streams that use modularity as the packaging technology. Therefore, no modular content is being distributed with RHEL 10.

Jira:RHELDOCS-20138[1]

10.4.4. Infrastructure services

Review functionalities deprecated for infrastructure services in Red Hat Enterprise Linux 10.0.

FTP clients and Servers software are now deprecated

The following FTP clients and servers software are deprecated and will be removed in the future major version of RHEL:

- **ftp**

- **lftp**
- **vsftpd**

These FTP protocol implementations are no longer under active development. We recommend that customers plan to migrate workflows based on FTP to one of either:

- OpenSSH and the **sftp** command, which provides an interactive interface for secure file transfer over the SSH protocol.
- WebDAV based on Apache httpd - various client implementations are available.

Jira:RHELDOCS-20610[1]

10.4.5. Networking

Review functionalities deprecated for networking in Red Hat Enterprise Linux 10.0.

ipset has been unmaintained

In RHEL 10, the **ipset** utility is unmaintained and is planned to be removed in a future major release. Red Hat will provide only critical bug fixes during the current release lifecycle. As an alternative to **ipset**, you can use the **nftables** sets functionality instead.

Jira:RHELDOCS-20147[1]

The BIND **auto-dnssec** parameter is deprecated

Starting with RHEL 9.7, the BIND **auto-dnssec** parameter is deprecated and will be removed in a future release. As a replacement, use the **dnssec-policy** parameter to specify a complete Key and Signing Policy (KASP) that groups all related configurations into a single, intuitive block.
For further details and information about migrating to **dnssec-policy**, see DNSSEC Key and Signing Policy in the BIND 9 upstream documentation.

Jira:RHELDOCS-21532[1]

10.4.6. File systems and storage

Review functionalities deprecated for file systems and storage in Red Hat Enterprise Linux 10.0.

The **squashfs** package has been deprecated

SquashFS is deprecated and will be removed in the next major release. It will no longer receive enhancements and is in RHEL 10 for specific use cases that are internal to Red Hat. Consider using EROFS as an alternative solution.

Jira:RHELDOCS-18450[1]

10.4.7. High availability and clusters

Review functionalities deprecated for high availability and clusters in Red Hat Enterprise Linux 10.0.

Deprecated High Availability Add-On features

The following features have been deprecated in Red Hat Enterprise Linux 10 and will be removed in the next major release:

- Specifying rules as multiple arguments. Use a single string argument instead.
- Specifying **score** as a standalone value in **pcs constraint location add** and **pcs constraint colocation ad**. Use **score=value** instead.
- Specifying the **--wait** option in resource commands except **pcs resource restart | move**, and in the commands **pcs cluster node add-guest | add-remote**. Use the following commands instead:
  - **pcs status wait** to wait for the cluster to settle into stable state.
  - **pcs status query resource** commands to verify that the resource is in the expected state after the wait.
- Using the **--force** flag to confirm potentially destructive actions such as **pcs cluster destroy**, **pcs quorum unblock**, **pcs stonith confirm**, **pcs stonith sbd device setup**, and **pcs stonith sbd watchdog test** commands. You should now use the **--yes** flag to confirm potentially destructive actions and reserve use of the **--force** flag to override validation errors.
- Using the **--force** flag to confirm overwriting files in **pcs cluster report**. Use the **--overwrite** flag instead.
- Assigning and unassigning ACL roles without specifying the **user** or **group** keyword.
- Configuring a score parameter in order constraints. The **pcs** command-line interface now produces a warning when a user attempts to configure a score parameter in order constraints.

Jira:RHELDOCS-19607[1]

10.4.8. Compilers and development tools

Review functionalities deprecated for compilers and development tools in Red Hat Enterprise Linux 10.0.

The **utmp** and **utmpx** interfaces in **glibc** are deprecated

The **utmp** and **utmpx** interfaces provided by the **glibc** library include a counter that counts time since the UNIX epoch. This counter will overflow on February 07, 2106. Therefore, **utmp** and **utmpx** are deprecated in RHEL 10 and will be removed in RHEL 11.

Jira:RHELDOCS-18080[1]

10.4.9. The web console

Review functionalities deprecated for the web console in Red Hat Enterprise Linux 10.0.

The host switcher in the RHEL web console is deprecated

The host switcher that provides connections to multiple machines through SSH from a single RHEL web console session is deprecated and disabled by default. Due to the web technology limitations, this feature cannot be secure.
In the short term, you can enable the host switcher after assessing the risks in your scenario with the **AllowMultiHost** option in the **cockpit.conf** file:

```
[WebService]
AllowMultiHost=yes
```

As more secure alternatives, you can use:

- the web console login page (with the secure limit of one host in a web browser session)
- the Cockpit Client flatpak

Jira:RHEL-4032[1]

10.4.10. Red Hat Enterprise Linux System Roles

Review functionalities deprecated for Red Hat Enterprise Linux system roles in Red Hat Enterprise Linux 10.0.

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-69315

The **sshd** variable deprecated and replaced by**sshd_config**

To unify coding standards across the RHEL system roles, the **sshd** variable has been replaced by the **sshd_config** variable. The **sshd** variable is now deprecated and might be removed from the **sshd** Ansible role in a future major version of RHEL.

Jira:RHEL-73440[1]

10.4.11. Virtualization

Review functionalities deprecated for virtualization in Red Hat Enterprise Linux 10.0.

libslirp has been deprecated

In RHEL 10, the **libslirp** networking back end has become deprecated, and will be removed in a future major version release.
Jira:RHEL-45147

The i440fx virtual machine type has been deprecated

In RHEL 10, the **i440fx** machine types for virtual machines (VMs) have become deprecated, and will be removed in a future major version of RHEL.
In addition, the **i440fx-rhel7.6** machine type has been replaced by **i440fx-rhel10.0**. As a consequence, a VM with a **i440fx-rhel7.6** machine type will not boot correctly after live migrating to a RHEL 10 host. Workaround: Restart the VM after live migration.

Jira:RHELDOCS-18672[1]

Legacy vCPU models are now deprecated

Several virtual CPU models are now deprecated and will become unsupported for use in virtual machines (VMs) in a future major release of RHEL. Notably, the deprecated models include the following:

- Intel Xeon 55xx and 75xx Processor families (also known as Nehalem)
- Intel Xeon v2 (also known as Ivy Bridge)
- AMD Opteron G4 and G5

To view the complete list of deprecated CPU models, use the following command:

```
# /usr/libexec/qemu-kvm -cpu help | grep depre | grep -v - -v
```

To check whether a running 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 'Nehalem'
```

Jira:RHEL-28971[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:RHELDOCS-20688[1]

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:RHELDOCS-20689[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:RHELDOCS-20691[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 10.2.

Jira:RHELDOCS-20692[1]

qcow2-v2 image format is deprecated

With RHEL 10.2, 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 10.2 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:RHELDOCS-20693[1]

10.4.12. Containers

Review functionalities deprecated for containers in Red Hat Enterprise Linux 10.0.

tzdata package is no longer installed by default in the minimal container images

The **tzdata** package is no longer installed in the **registry.access.redhat.com/ubi10-minimal** container image. As a consequence, if you migrate your minimal container builds from a previous RHEL release to RHEL 10.0, and you enter the **microdnf reinstall tzdata** command to reinstall the **tzdata** package, you get an error message because the **tzdata** package is no longer installed by default. In this case, enter the **microdnf install tzdata** command to install **tzdata**.

Jira:RHELDOCS-18700[1]

The Podman v5.0 deprecations

In RHEL 10.0, 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 **containernetworking-plugins** package and the CNI network stack are no longer supported.

- If you upgrade from the previous RHEL versions to RHEL 10.0 or if you have a fresh installation of RHEL 10.0, the CNI is no longer available. As a result, you have to run the **podman rmi --all --force** command to remove all images and containers that are using those images.
- If present, the **cni** value in the containers.conf file for the **network_backend** option must be changed to **netavark** or can be unset.

Jira:RHEL-40641

The **podman-tests** package has been deprecated

The **podman-tests** package has been deprecated in the AppStream repository. The package is now available in the CodeReady Linux Builder (CRB). More information about the CRB repository can be found at The CodeReady Linux Builder repository .
Jira:RHEL-67860

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]

10.5. DEPRECATED FEATURES IDENTIFIED IN PREVIOUS RELEASES

Review functionalities that are deprecated in earlier Red Hat Enterprise Linux versions.

10.5.1. SSSD

Review functionalities deprecated for SSSD in previous Red Hat Enterprise Linux versions.

The SMB1 protocol is deprecated in Samba

Starting with Samba 4.11, the insecure Server Message Block version 1 (SMB1) protocol is deprecated and will be removed in a future release.
To improve the security, by default, SMB1 is disabled in the Samba server and client utilities.

Jira:RHELDOCS-16612[1]

10.6. DEPRECATED PACKAGES

Review the packages that are deprecated in Red Hat Enterprise Linux 10.2. Although these packages remain fully supported in this release, they will likely be removed in a future major version.

IMPORTANT

The support status of deprecated packages remains unchanged within RHEL 10.

The following packages have been deprecated in RHEL 10:

- daxio

- ftp
- gvisor-tap-vsock-gvforwarder
- lftp
- libpmem
- libpmem2
- libpmemblk
- libpmemlog
- libpmemobj
- libpmemobj-cpp
- libpmempool
- libslirp
- nvml
- pmempool
- pmreorder
- sdl2-compat
- vsftpd
- wget

CHAPTER 11. KNOWN ISSUES

Understand how newly identified and previously known issues might affect your use of Red Hat Enterprise Linux 10.2, and how to work around them.

A known issue is listed in all future release notes until resolved, at which point it is published as a fixed issue.

11.1. INSTALLER AND IMAGE CREATION

Review known issues for installer and image creation in Red Hat Enterprise Linux 10.2.

Boot container installation in UEFI mode fails on systems without **/boot**

The installation of a bootable container image in UEFI mode fails in the absence of the separate **/boot** partition.
Workaround: Ensure the Kickstart file or manual configuration includes both an EFI System partition and a distinct **/boot** partition in addition to the root (/) partition.

Jira:RHEL-147437

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-80672

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-58827

11.2. SECURITY

Review known issues for security in Red Hat Enterprise Linux 10.2.

rust-rpm-sequoia fails when importing OpenPGP certificates with keys disallowed by **crypto-policies**

Importing OpenPGP certificates that contain keys disallowed by the system-wide cryptographic policy causes the **rust-rpm-sequoia** library to fail. Consequently, the failure of the import prevents further operations, such as importing additional OpenPGP certificates from a single file.

To work around this problem, remove the disallowed key from the file before importing the certificate bundle. As a result, **rust-rpm-sequoia** does not fail when you import OpenPGP certificates only with keys allowed by **crypto-policies**.

Jira:RHEL-152461

11.3. RHEL FOR EDGE

Review known issues for RHEL for Edge in Red Hat Enterprise Linux 10.2.

Greenboot triggers a warning message during the first boot

When booting a system for the first time with the **greenboot-0.16.2-0** package, the system might log a **WARN** message stating that boot data is unavailable. This occurs because the initial boot data has not yet been generated. This message is benign and does not affect the system's operation or the health check process. You can safely ignore the warning during the initial boot.
Jira:RHEL-141567

11.4. SOFTWARE MANAGEMENT

Review known issues for software management in Red Hat Enterprise Linux 10.2.

DNF installs a package from a local file when the package version is excluded in versionlock

When you exclude a package version in the **versionlock** DNF plugin configuration, DNF still installs the specified package version from a package local file.
To work around this problem, complete the following steps:

1. Turn a directory with local packages into a local repository by using the **createrepo_c** tool.
2. Enable the local repository in the DNF configuration.
3. Install all packages by their names.

As a result, the **versionlock** plugin applies to packages from the local repository and has no effect on directory with local package files.

NOTE

Consider not installing packages by a local file path if you do not want certain package versions to be installed.

For more information, see the **dnf-versionlock(8)** man page on your system.

Jira:RHEL-94828

11.5. NETWORKING

Review known issues for networking in Red Hat Enterprise Linux 10.2.

RHEL does not contain closed-source modem unlocking tools

Federal Communications Commission (FCC) regulations require that modems in the United States

must be enabled by using an unlocking tool from the modem manufacturer. RHEL does not provide these tools if they are closed-source software according to FCC regulations. However, they might be available in an unsupported third-party repository, such as RPM Fusion.
For further details, see Installing the FCC unlocking tool for modems from third-party repositories .

Jira:RHEL-100066[1]

Preventing non-root users from creating system-wide NetworkManager connection profiles

You can set certain properties in NetworkManager connection profiles, such as **802-1x.client-cert**, to a path to a certificate file. Because the **NetworkManager** service runs as the **root** user, the service can access those files independent of their file permissions. This can lead to security problems in the following scenarios:

- A user creates a private connection profile and specifies a path to another user's certificate file.
  With NetworkManager in RHEL 10.2 and later, referring to other users' certificates in private profiles is no longer possible.

- A user creates a system-wide connection profile and specifies a path to another user's certificate.
  On RHEL, users can only create system-wide profiles if they are logged in locally to the console and not remotely, such as over SSH. To not change this behavior of NetworkManager during the RHEL 10 release cycle, users can still create system-wide profiles.

  To mitigate the risk, you can prevent normal users from creating system-wide connection profiles. For example, create the **/etc/polkit-1/rules.d/20-nm-non-root.rules** file with the following content:

```
polkit.addRule(function(action, subject) {
    if (action.id == "org.freedesktop.NetworkManager.settings.modify.system" &&
        !subject.isInGroup("wheel")) {
        return polkit.Result.AUTH_ADMIN_KEEP;
    }
});
```

  The setting takes effect immediately.

Jira:RHELDOCS-21618[1]

11.6. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Review known issues for dynamic programming languages, web and database servers in Red Hat Enterprise Linux 10.2.

SELinux prevents replication in a Galera cluster after upgrading from MariaDB 10.11 to MariaDB 11.8

If you run a Galera cluster and upgrade MariaDB 10.11 to MariaDB 11.8, SELinux prevents replication among the cluster nodes.
There is no known workaround for the issue, while keeping SELinux in the **enforcing** mode.

Note that setting the SELinux mode to **permissive** is not a secure workaround, because this significantly reduces the security of your servers.

Jira:RHEL-171580

11.7. IDENTITY MANAGEMENT

Review known issues for Identity Management (IdM) in Red Hat Enterprise Linux 10.2.

ipa-migrate does not migrate SSH public keys

When migrating an Identity Management (IdM) deployment using the **ipa-migrate** tool, SSH public keys assigned to user accounts and ID overrides are not transferred to the destination server. As a consequence, users cannot authenticate using SSH public key authentication after migration.
To work around this problem, retrieve the SSH public keys from the source server using the **ipa user-find --all** or **ldapsearch** commands, and then re-add them on the destination server using the **ipa user-mod --sshpubkey** command.

RHEL-147173

Jira:RHEL-147173

11.8. VIRTUALIZATION

Review known issues for virtualization in Red Hat Enterprise Linux 10.2.

High-memory Windows guests might fail to validate with SVVP

Currently, when using the Server Virtualization Validation Program (SVVP) software to validate a Windows virtual machine (VM) with a large amount of assigned memory, the validation might fail with a **GetPhysicallyInstalledSystemMemory failed** error message. As a consequence, the VM cannot be validated for SVVP support.
Jira:RHEL-81999

VMs on IBM Z hosts sometimes fail to boot when an invalid boot device is specified

Currently, for KVM virtual machines (VMs) hosted on IBM Z systems, the boot order setting does not work consistently. If the boot device configured with **<boot order='1'/>** is invalid, the next boot device specified by the boot order setting sometimes does not boot correctly. This might cause the VM to shut down or become unresponsive.

Jira:RHEL-151317[1]

Stop errors in Windows guests

Currently, in virtual machines that use Windows guest operating systems on RHEL hosts, a variety of stop errors (also known as BSOD) might occur. For details of the known errors, see List of known Windows BSOD issues on OpenShift Virtualization and RHEL KVM on Red Hat Knowledge Base. For instructions on troubleshooting the errors, see Recommendations when investigating Windows BSOD issues.

Jira:RHELDOCS-22157[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

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]

11.9. RHEL IN CLOUD ENVIRONMENTS

Review known issues for RHEL in cloud environments in Red Hat Enterprise Linux 10.2.

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-75576[1]

11.10. CONTAINERS

Review known issues for containers in Red Hat Enterprise Linux 10.2.

EUS repositories are not enabled by default

RHEL 10 systems target the major release version, such as RHEL 10. Standard repositories remain active by default, preventing the automatic enablement and correct path resolution for Extended Update Support (EUS) repositories. Consequently, systems can not receive the expected EUS-specific software updates.
If the release version is not set, DNF attempts to pull metadata from a major-version URL such as **.../rhel10/10/...**, which does not exist for EUS, resulting in a **404 Not Found** error during the metadata sync.

To work around this problem, follow these steps:

- Override the default DNF variable to point to the specific EUS minor version, such as **10.2**. This ensures the repository URL construction is accurate.

```
# Set releasever to EUS version (mandatory)
RUN echo "10.0" > /etc/dnf/vars/releasever
```

After setting **releasever**, use one of the following two approaches to enable EUS:

  - Command-line configuration:
    - Manually disable the standard repositories and enable the EUS versions.

```
# dnf config-manager --set-disabled <standard_repo_id>
# dnf config-manager --set-enabled <eus_repo_id>
```

  - Manual File editing:
    1. Run the **dnf repolist** command to trigger the **rhsm dnf** plugin to generate **/etc/yum.repos.d/redhat.repo** dynamically.
    2. Edit the relevant **.repo** files in the **/etc/yum.repos.d/** directory to set **enabled=0** for standard repositories and **enabled=1** for EUS repositories.

Jira:RHELDOCS-21758[1]

11.11. KNOWN ISSUES IDENTIFIED IN RHEL 10.1

Review known issues identified in Red Hat Enterprise Linux 10.1 that might affect your use of Red Hat Enterprise Linux 10.2.

11.11.1. Installer and image creation

Review known issues for installer and image creation identified in Red Hat Enterprise Linux 10.1.

Crash dumps are not performed by default

By default, crash dumps do not occur for default installation methods using RHEL Image Mode, because the **crashkernel=** kernel argument is not set. To work around this problem, set a **crashkernel=** kernel argument at build or during installation time.
Jira:RHEL-82380

11.11.2. Security

Review known issues for security identified in Red Hat Enterprise Linux 10.1.

Containers fail to start when **fapolicyd** is running

The **fapolicyd** framework does not fully support namespaces and containers. As a consequence, containers fail to start when **fapolicyd** is running.
To work around this problem, create the **/etc/fapolicyd/rules.d/25-runc.rules** file with the following content:

```
allow perm=any pattern=ld_so exe=/usr/bin/runc : all
allow perm=any uid=0 pattern=ld_so exe=/runc : trust=1
```

Save the file, run the **fagenrules** script, and enter the **fapolicyd-cli --reload-rules** command to apply the changes. Alternatively, you can remove the **tmpfs** value from the **watch_fs** option in the **/etc/fapolicyd/fapolicyd.conf** file and restart the **fapolicyd** service by using the **systemctl restart fapolicyd** command, but this lowers the system security.

As a result, you can use **fapolicyd** on systems running containers after you apply the previously described workaround. This preserves the enhanced security provided by **fapolicyd** and helps comply with configuration standards such as the Security Technical Implementation Guide (STIG) from the Defense Information Systems Agency (DISA).

Jira:RHEL-114562

sq cannot generate keys in FIPS mode

The **sq** utility from the Sequoia PGP toolset uses the deprecated OpenSSL API for key generation. Consequently, you cannot generate keys with **sq** on the system running in FIPS mode.
Jira:RHEL-85985

GnuTLS cannot convert ML-DSA private keys to public ones

GnuTLS lacks an algorithm to convert a private ML-DSA key in the expanded form to a public ML-DSA key. Consequently, operations requiring both keys fail when only the expanded private key is provided.
Workaround: Use the **openssl** command to convert such a private key to a public key: **openssl dsa -in *<private_key>* -pubout -out *<public_key>***. As a result, the public key is available for use in other operations.

Jira:RHEL-102992

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 hard-coded on the **crypto-policies** level. As a result, PQ algorithms for **rpm-sequoia** are enabled regardless of any settings in **crypto-policies**.

Jira:RHEL-112392

11.11.3. Shells and command-line tools

Review known issues for shells and command-line tools identified in Red Hat Enterprise Linux 10.1.

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

11.11.4. Networking

Review known issues for networking identified in Red Hat Enterprise Linux 10.1.

Inbound IPsec cryptographic offload can fail in SR-IOV switchdev mode with SMFS

If you configure IPsec cryptographic offload on a Mellanox ConnectX network interface controller (NIC) in Single-Root I/O Virtualization (SR-IOV) **switchdev** mode with the flow steering mode set to Software Managed Flow Steering (SMFS), the hardware offload for inbound IPsec Security Associations (SAs) fails. In this case, the **ip xfrm state dir in show** command returns the following error:

```
Error: mlx5_core: Device failed to offload this state.
```

To work around this problem, switch to Device-Managed Flow Steering (DMFS) before switching the device to **switchdev** mode. By using DMFS, the inbound IPsec state can successfully be offloaded to the hardware.

Jira:RHEL-114861[1]

11.11.5. File systems and storage

Review known issues for file systems and storage identified in Red Hat Enterprise Linux 10.1.

iSCSI-backed logical volumes fail to activate after a reboot

During installation, a logical volume spanning a local disk and an iSCSI device can fail to activate the iSCSI device in the installed system. This occurs where a non-root filesystem LVM logical volume is located both on a local disk and on an iSCSI device, which results in the iSCSI device not getting configured with **node.startup=onboot** by the installation program. As a result, the system cannot access the volume after reboot, because it doesn't get automatically activated upon boot.
Workaround: Manually create the logical volume after the installation or update the iSCSI node configuration by setting **node.startup=automatic** in the relevant file in the **/var/lib/iscsi/nodes/** directory.

Jira:RHEL-53719

11.11.6. Dynamic programming languages, web and database servers

Review known issues for dynamic programming languages, web and database servers identified in Red Hat Enterprise Linux 10.1.

MySQL does not work with RHEL in image mode

The MySQL database management systems in RHEL 10 do not use the **sysusers.d** directories to populate users and working directories. Additionally, MySQL also does not use the **tmpfiles.d** directory. As a consequence, the database user can be missing and MySQL is not able to initialize because its working directory is missing. There is currently no workaround for this issue.

Jira:RHELDOCS-21374[1]

11.11.7. Desktop

Review known issues for desktop identified in Red Hat Enterprise Linux 10.1.

Plymouth duplicates log entries of the kernel log ring buffer

Plymouth, an application which provides a graphical boot experience for Red Hat Enterprise Linux, has a "console syndication" feature that outputs log messages to all configured consoles during boot. The kernel can natively output log messages only to the last configured console. In the default configuration, the kernel is muted, but removing the **quiet** argument from the kernel command line unmutes the kernel, and causes both Plymouth and the kernel to send the boot log messages to the last-configured console. As a result, log messages might be duplicated on the last-configured console (for example ttyS0). Plymouth further duplicates these log entries by replaying the entire contents of the kernel log ring buffer during boot and shutdown. To work around this problem, disable Plymouth.

Jira:RHEL-60198[1]

11.11.8. Red Hat Enterprise Linux System Roles

Review known issues for Red Hat Enterprise Linux system roles identified in Red Hat Enterprise Linux 10.1.

Ansible **rpm_key** modules fail to work with the OpenPGP v6**RPM-GPG-KEY-redhat-release** key

RHEL 10.1 uses the Red Hat RPM signing key extended with a post-quantum public key and stored in the **/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release** file in the OpenPGP v6 format. Because the Ansible **rpm_key** modules use the GnuPG tools, which cannot handle post-quantum keys and OpenPGP v6, the modules fail to work with this key.
Jira:RHEL-126844

11.11.9. Virtualization

Review known issues for virtualization identified in Red Hat Enterprise Linux 10.1.

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]

VMs with SEV-SNP enabled fail to boot with arch-capabilities=on

Virtual machines (VMs) with SEV-SNP enabled fail to boot when using the **arch-capabilities=on** CPU flag.
To work around this problem, disable the **arch-capabilities** feature in the CPU section of the VM's configuration:

```
<cpu mode='host-passthrough' check='none'>
  <feature name='arch-capabilities' policy='disable'/>
</cpu>
```

Jira:RHEL-100313[1]

11.11.10. RHEL Lightspeed

Review known issues for RHEL Lightspeed identified in Red Hat Enterprise Linux 10.1.

The command-line assistant cannot verify the Satellite server certificate

The command-line assistant does not recognize the Satellite certificate authority (CA) certificate for the Red Hat Satellite server. The Satellite CA certificate is used to issue and sign certificates for hosts that register with and are managed by Satellite. As a consequence, the command-line assistant cannot establish secure connections to the Satellite server, which prevents it from functioning correctly.
Work around: copy the Satellite CA certificate to the system trust store and update the CA trust database:

```
$ sudo cp /etc/rhsm/ca/katello* /etc/pki/ca-trust/source/anchors/
$ sudo update-ca-trust
```

Jira:RHELDOCS-21325[1]

Security risk when using unsupported AI models

Do not use unsupported models. Changing the RHEL Offline Container model to unsupported models might allow the execution of arbitrary code or compromise the integrity of Red Hat Enterprise Linux (RHEL).
No known workaround exists.

Jira:RHELDOCS-21726[1]

11.12. KNOWN ISSUES IDENTIFIED IN RHEL 10.0

Review known issues identified in Red Hat Enterprise Linux 10.0 that might affect your use of Red Hat Enterprise Linux 10.2.

11.12.1. Installer and image creation

Review known issues for installer and image creation identified in Red Hat Enterprise Linux 10.0.

Podman and bootc do not share the same registry login process

Podman and **bootc** use different registry login processes when pulling images. As a consequence, if you login to an image by using Podman, logging to a registry for **bootc** will not work on that image. When you install an image mode for RHEL system, and login to registry.redhat.io by using the following command:

```
# podman login registry.redhat.io <username_password>
```

And then you attempt to switch to the **registry.redhat.io/rhel9/rhel-bootc** image with the following command:

```
# bootc switch registry.redhat.io/rhel9/rhel-bootc:9.4
```

You should be able to see the following message:

```
Queued for next boot: registry.redhat.io/rhel9/rhel-bootc:9.4
```

However, an error is displayed:

```
ERROR Switching: Pulling: Creating importer: Failed to invoke skopeo proxy method OpenImage: remote error: unable to retrieve auth token: invalid username/password: unauthorized: Please login to the Red Hat Registry using your Customer Portal credentials. Further instructions can be found here: https://access.redhat.com/RegistryAuthentication
```

Workaround: Follow the steps Configuring container pull secrets to use authenticated registries with **bootc**.

Jira:RHELDOCS-18471[1]

cloud-init growpart skips with composefs is enabled

When composefs is enabled, if you generate an image from the generic base image, then the rootfs will not grow the filesystem, prompting an error similar to:

```
2024-04-30 17:27:53,543 - cc_growpart.py[DEBUG]: '/' SKIPPED: stat of 'overlay' failed: [Errno 2]
No such file or directory: 'overlay'
```

Workaround: You can add a custom growpart, by specifying the **rootfs** default size in the container, instead of dynamically choosing 100G at instance creation time to be able to write a partitioning config in the container.

Jira:RHEL-34859

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

The installation program becomes unresponsive during final RPM installation stage

An installation program might become unresponsive during the RPM installation process at the final stage. Before the issue occurs, you might see the repeated **Configuring rootfiles.noarch** messages.
Workaround: Restart the installation process.

Jira:RHEL-67865[1]

Disabled keyboard layout switching by using shortcut during installation

To prevent confusion caused by a broken keyboard shortcut to change keyboard layout, this feature has been disabled in Anaconda. You cannot change keyboard layouts by using shortcuts during installation. Workaround: Use the keyboard layout icon on the top bar to switch layouts.
Jira:RHEL-74504

The installation program now respects the **BOOTIF** boot argument

Previously, the RHEL installation program 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-69400[1]

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 through 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-19853[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-83577

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-58829

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 might 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 might not function correctly on s390x and ppc64le architectures.

Jira:RHELDOCS-19496[1]

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]

Installation fails due to busy partitions

A race condition in the storage subsystem causes the installation to fail when writing the partition table to disk. The system displays the following error message:

> Partition(s) have been written, but we have been unable to inform the kernel of the change.

This error occurs because the partitions are reported as busy and the changes cannot be synchronized. To work around this problem, restart the installation.

Jira:RHEL-158237

11.12.2. Security

Review known issues for security identified in Red Hat Enterprise Linux 10.0.

SELinux policy rules for four**libvirt** services temporarily changed into permissive mode

Previously, the SELinux policy was changed to reflect the replacement of the legacy monolithic **libvirtd** daemon with a new set of modular daemons. Because this change requires testing of many scenarios, the following services have been temporarily changed into SELinux permissive mode:

- **virtqemud**
- **virtvboxd**
- **virtstoraged**
- **virtsecretd**

To prevent harmless AVC denials, **dontaudit** rules have been added to the SELinux policy for these services.

Jira:RHEL-77808[1]

Cryptographic tokens do not work in FIPS mode with**pkcs11-provider**

When the system runs in FIPS mode, the **pkcs11-provider** OpenSSL provider does not work correctly and the OpenSSL TLS toolkit falls back to the default provider. Consequently, OpenSSL fails to load PKCS #11 keys, and cryptographic tokens do not work in this scenario.

To work around this problem, set the **pkcs11-module-assume-fips = true** parameter in the PKCS #11 section of the **openssl.cnf** file. See the **pkcs11-provider(7)** man page on your system for more information. With this configuration change, **pkcs11-provider** works in FIPS mode.

Jira:RHEL-68621

OpenSSL stores ML-KEM and ML-DSA private keys in standard formats

In RHEL 10.0, the open quantum-safe provider for OpenSSL (**oqsprovider**) generated private keys in a format that did not conform to any of the file formats proposed by the IETF LAMPS work group. Consequently, the key files were unreadable by other applications that follow the IETF standard and could not be handled by applications that require providing the key in the seed format for import. With this update, OpenSSL no longer uses **oqsprovider** and its post-quantum cryptography (PQC) implementation generates the keys in standard formats. As a result, you can use OpenSSL ML-KEM and ML-DSA keys for storing long-term secrets.
Jira:RHEL-72719

11.12.3. Shells and command-line tools

Review known issues for shells and command-line tools identified in Red Hat Enterprise Linux 10.0.

uname command produces an unknown output

The **uname** command displays unknown output with flags **--hardware-platform** and **--processor**. In the previous RHEL versions, **uname -i** and **uname -p** were aliases for **uname -m** and are not portable even across GNU or Linux distributions.
As a workaround, you can use the **-m** flag instead of the **-i** and **-p** flags.

Jira:RHEL-74146

11.12.4. Infrastructure services

Review known issues for infrastructure services identified in Red Hat Enterprise Linux 10.0.

Nginx does not support PKCS #11 and TPM

The OpenSSL engines API was deprecated in RHEL 9 and removed from Nginx in RHEL 10. The corresponding functionality using the current OpenSSL providers API is not yet available. As a consequence, the Nginx HTTP server does not work with hardware security modules (HSMs) through PKCS #11 and Trusted Platform Module (TPM) devices.
Jira:RHEL-33742

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-19770[1]

11.12.5. Networking

Review known issues for networking identified in Red Hat Enterprise Linux 10.0.

VMware vCenter now correctly removes a SATA disk from a running RHEL VM

When using the VMware vCenter interface to remove a SATA disk from a running RHEL 10 guest on the VMware ESXi hypervisor, the disk previously 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. This update fixes the issue, and the SATA disk is fully removed in the described scenario.

Jira:RHEL-79913[1]

11.12.6. High availability and clusters

Review known issues for high availability and clusters identified in Red Hat Enterprise Linux 10.0.

ACL roles should not reference location constraints with two rules

In Red Hat Enterprise Linux 10, more than one top-level rule in a location constraint is not supported. When upgrading from RHEL 9 to RHEL 10, verify that any ACL roles you have configured do not reference a location constraint with two rules and are still valid.
Jira:RHEL-62722

11.12.7. Compilers and development tools

Review known issues for compilers and development tools identified in Red Hat Enterprise Linux 10.0.

The new version of TBB is incompatible

RHEL 10 includes the Threading Building Blocks (TBB) library version 2021.11.0, which is incompatible with the versions distributed with previous releases of RHEL. You must rebuild applications that use TBB to make them run on RHEL 10.
Jira:RHEL-33633

11.12.8. Identity Management

Review known issues for Identity Management (IdM) identified in Red Hat Enterprise Linux 10.0.

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]

Installing a RHEL 7 IdM client with a RHEL 10 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 10 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 10 fails.
Workaround: Upgrade the host to RHEL 8 or later before installing an IdM client on it.

Jira:RHELDOCS-19015[1]

ldapmodify does not delete a single specific value from any attribute in **cn=config**

Currently, when you try to delete a value from any attribute in **cn=config**, the value remains in the attribute and the server might require a restart to fully remove it.
Workaround: Remove the entire attribute, including all its values, by performing a modify operation without specifying any values. Then re-add the values you need. Alternatively, use the following **dsconf** command to remove a specific value without a server restart:

```
# dsconf <instance_name> config delete <attribute_name>=<undesired_value>
```

Jira:RHEL-25071

11.12.9. SSSD

Review known issues for SSSD identified in Red Hat Enterprise Linux 10.0.

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]

11.12.10. Desktop

Review known issues for desktop identified in Red Hat Enterprise Linux 10.0.

Standard mouse cursor is offset in VMs when using Mutter

When you use a standard mouse within a virtual machine (VM) configuration in the Mutter compositing window manager, you might notice an offset between the physical mouse cursor and the

actual pointer within the virtual environment. The actual pointer might not even be visible in the virtual environment.
Workaround: If your scenario requires precise input, use a tablet as an input device in the VM configuration.

Jira:RHEL-69291

11.12.11. Graphics infrastructures

Review known issues for graphics infrastructures identified in Red Hat Enterprise Linux 10.0.

Standard mouse cursor is offset in VMs when using Mutter

When you use a standard mouse within a virtual machine (VM) configuration in the Mutter compositing window manager, you might notice an offset between the physical mouse cursor and the actual pointer within the virtual environment. The actual pointer might not even be visible in the virtual environment.
Workaround: If your scenario requires precise input, use a tablet as an input device in the VM configuration.

Jira:RHEL-45898

11.12.12. The web console

Review known issues for the web console identified in Red Hat Enterprise Linux 10.0.

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]

11.12.13. Red Hat Enterprise Linux System Roles

Review known issues for Red Hat Enterprise Linux system roles identified in Red Hat Enterprise Linux 10.0.

ansible-core does not install **sshpass** as a dependency

The **ansible-core** package does not install the **sshpass** package as a dependency. Consequently, you cannot use Ansible to manage systems over SSH with an SSH password.
Workaround: On the control node, manually install **sshpass** after you install **ansible-core**. As a result, you can use Ansible in the scenario described above.

Jira:RHEL-86829[1]

11.12.14. Virtualization

Review known issues for virtualization identified in Red Hat Enterprise Linux 10.0.

QEMU no longer prevents using SEV-SNP

Previously, when attempting to start a virtual machine (VM) with AMD SEV-SNP enabled, QEMU checked the incorrect capability of KVM, and the guest failed to start. As a consequence, running VMs with AMD SEV-SNP configured was not possible with RHEL10. This problem has been fixed, and running VMs with SEV-SNP works as expected now.

Jira:RHEL-58928[1]

Network boot for VMs now works correctly without an RNG device

Previously, when a virtual machine (VM) did not have an RNG device configured and its CPU model did not support the RDRAND feature, it was not possible to boot the VM from the network. With this update, the problem has been fixed, and VMs that do not support RDRAND can boot from the network even without an RNG device configured.

Note, however, that adding an RNG device is highly encouraged for VMs that use a CPU model that does not support RDRAND, in order to increase security when booting from the network.

Jira:RHEL-66234

RHEL 10 guests no longer crash on restart in Google Cloud and Alibaba

When using a RHEL 10.0 instance on Google Cloud or the Alibaba Cloud, restarting the instance previously caused a kernel panic in the guest operating system if the **virtio-net** driver was in use. This issue has been fixed and RHEL 10 guests no longer crash in the described scenario.

Jira:RHEL-56981[1]

Secure Execution VMs can now boot with file-backed memory backing

Previously, if you configured a virtual machine (VM) with enabled Secure Execution to use file-backed memory backing, the VM failed to boot, and instead displayed a **Protected boot has failed** error. Now, the VM boots as expected.

Jira:RHEL-58218

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

VMs with large memory can now boot correctly on SEV-SNP host with AMD Genoa CPUs

Previously, virtual machines (VMs) could not boot on hosts that used a 4th Generation AMD EPYC processor (also known as Genoa) and had the AMD Secure Encrypted Virtualization with Secure Nested Paging (SEV-SNP) feature enabled. Instead of booting, a kernel panic occurred in the VM. This issue has now been fixed.

Jira:RHEL-32892[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

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]

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

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-82685

Enabling 3D support no longer prevents installing a RHEL 10 guest on ESXi

Prior to this update, if you selected the Enable 3D support option in VMware ESXI for installing a RHEL 10 guest operating system, the installation did not start correctly, and instead showed a blank screen. This issue has been fixed, and you can now install RHEL 10 guests in the described scenario.

Jira:RHEL-88668[1]

11.12.15. RHEL in cloud environments

Review known issues for RHEL in cloud environments identified in Red Hat Enterprise Linux 10.0.

RDMA devices currently do not work on vSphere

When using a RHEL 10 instance on the VMware vSphere platform, the **vmw_pvrdma** module currently does not install properly. As a consequence, VMware paravirtual remote direct memory access (PVRDMA) devices do not work on the affected instances.

Jira:RHEL-41133[1]

The leapp upgrade fails when upgrading from RHEL 9.6 to RHEL 10.0 for the cloud-init network configuration

If you deploy RHEL 9.6 with the **cloud-init** default configuration and with **sysconfig** as the default network configuration directory, the **sysconfig** configuration files do not support the **ifcfg** legacy format for RHEL 10.0. Consequently, the **leapp** upgrade fails when upgrading from RHEL 9.6 to RHEL 10.0 for the legacy network configuration files, such as ifcfg-*<enp1s0>*.
Workaround: Convert the **sysconfig** configuration files into the NetworkManager native **keyfile** format:

1. Modify the connection:

```
# nmcli connection modify "System <enp1s0>" connection.id "cloud-init <enp1s0>"
```

2. Migrate the connection:

```
# nmcli connection migrate /etc/sysconfig/network-scripts/ifcfg-<enp1s0>
```

3. Move the connection profile:

```
# sudo mv /etc/NetworkManager/system-connections/"cloud-init <enp1s0>.nmconnection"
/etc/NetworkManager/system-connections/cloud-init-<enp1s0>.nmconnection
```

4. Reload the network connection settings:

```
# nmcli conn reload
```

As a result, the leapp upgrade from RHEL 9.6 to RHEL 10.0 now works with the updated configuration.

Jira:RHEL-82209[1]

Upgrading a RHEL 9.6 guest on VMware ESXi to RHEL 10.0 causes **cloud-init** to rewrite the network configuration

After a upgrading a RHEL guest on the VMware ESXi hypervisor from RHEL 9.6 to RHEL 10.0, the **cloud-init** tool currently cannot detect the VMware data source and cannot restore its configuration from the cache. As a consequence, **cloud-init** reverts to the **None** data source, and rewrites the network configuration of the guest.
Workaround: Remove the **disable_vmware_customization** flag from the **/etc/cloud/cloud.cfg** file before you reboot the guest during the upgrade process. As a result, the upgraded guest will retain its previous network configuration.

Jira:RHEL-82210[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]

11.12.16. Containers

Review known issues for containers identified in Red Hat Enterprise Linux 10.0.

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

11.12.17. RHEL Lightspeed

Review known issues for RHEL Lightspeed identified in Red Hat Enterprise Linux 10.0.

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]

11.13. KNOWN ISSUES IDENTIFIED IN PREVIOUS RELEASES

Review known issues identified in earlier Red Hat Enterprise Linux versions that might affect your use of Red Hat Enterprise Linux 10.2.

11.13.1. Networking

Review known issues for networking identified in previous Red Hat Enterprise Linux versions.

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:RHELDOCS-20686[1]

kTLS does not support offloading of TLS 1.3 to NICs

Kernel Transport Layer Security (kTLS) does not support offloading of TLS 1.3 to NICs. Consequently, software encryption is used with TLS 1.3 even when the NICs support TLS offload.
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:RHELDOCS-20687[1]

11.13.2. Virtualization

Review known issues for virtualization identified in previous Red Hat Enterprise Linux versions.

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]

CHAPTER 12. FIXED ISSUES

Review issues that have been fixed in Red Hat Enterprise Linux 10.2.

12.1. INSTALLER AND IMAGE CREATION

Review issues that have been fixed for installer and image creation in Red Hat Enterprise Linux 10.2.

Root passwords are correctly set in ISO images

Before this update, the blueprint incorrectly used the **user** command to configure the **root** password for ISO images. As a consequence, the root password was not set, preventing users from logging in. With this release, the root password is correctly applied for the root user during the ISO image build process, enabling successful system access.
Jira:RHEL-4644

Installer falls back to English in text mode for unsupported languages

Before this release, the installer did not set the display mode (text, graphical, or non-interactive) early enough during startup. As a result, the check to determine whether a selected language is supported in text mode did not run. In text mode installations, languages that are not supported in the text user interface, such as Japanese, could be used, resulting in unreadable output.
With this fix, the installer correctly detects languages that are not supported in the text mode. If an unsupported language is selected, the text user interface falls back to English. The installed system is still configured to use the originally selected language.

Jira:RHEL-16168

The driver disk menu now correctly displays user input on the console

Before this release, when starting a RHEL installation with the **inst.dd** kernel command-line option, the console failed to render characters typed by the user. As a consequence, the lack of visual feedback made the application appear unresponsive, even though the input was still being processed in the background. With this update, this display issue has been resolved, and user input is now visible as expected during the driver disk selection process.
Jira:RHEL-58828

Anaconda installation program no longer fails if/**boot** partition is not created when using the **ostreecontainer** Kickstart command

Before this update, while using the **ostreecontainer** Kickstart command to install a bootable container, the /**boot** partition was not created. As a consequence, the installation failed because it required a dedicated /**boot** partition to proceed with the container deployment.
With this update, you can use Anaconda to install a bootable container image without having a separate /**boot** partition.

Jira:RHEL-66155

Rescue mode in Anaconda now correctly detects and mounts image-based systems

Before this release, the rescue scanner failed to identify image-based installations due to their unique filesystem hierarchy. The environment now automatically mounts the system under the /**mnt**/**sysroot** mount point and provides the specific **chroot** command required to access the active

deployment. As the image-based systems are immutable, manual changes should be limited to **/etc** or **/var**.
Jira:RHEL-135116

12.2. SECURITY

Review issues that have been fixed for security in Red Hat Enterprise Linux 10.2.

AIDE no longer terminates when a monitored file is changed

Before this update, AIDE terminated with an error if a file was truncated or removed while AIDE was computing its hash. With this update, AIDE detects when a file is truncated or deleted during hash calculation and handles the condition safely. As a result, AIDE successfully completes integrity checks even if a monitored file change size or is removed during processing.
Jira:RHEL-1383

fapolicyd-cli --check-trustdb no longer reports files without size or checksum information

Some files, for example, **/usr/lib/rpm/redhat/redhat-annobin-cc1** or **/etc/selinux/targeted/policy/policy.33**, owned by an RPM package, are expected to be changed during and after the installation, but they are still owned by the corresponding package. Consequently, **fapolicyd** cannot verify such files. With this release, the **fapolicyd** framework no longer adds files that do not have size or checksum information in the RPM database to the trust database. As a result, the **fapolicyd-cli --check-trustdb** command does not report the **miscompares: size sha256** error message for such files.
Jira:RHEL-94786

Remote serving of PKCS #11 tokens is no longer broken in **p11-kit**

Before this update of the **p11-kit** packages, a zero-length recursive attribute array was improperly read in the remote procedure call (RPC) mechanism. Consequently, remote serving of PKCS #11 tokens broke due to a communication error. This update fixes the reading of zero-length attribute arrays. As a result, a **p11-kit** server can remotely serve PKCS #11 tokens.

Jira:RHEL-97770[1]

NSS database password updates no longer corrupt ML-DSA seeds

Before this update, a bug in how NSS handled database re-encryption prevented the ML-DSA seed attribute from updating when you changed the database password. As a result, the seed value was permanently lost, even if you knew the previous password.
With this update, password changes correctly update the ML-DSA seed attribute and no longer cause the permanent loss of seed values. Previously lost seeds cannot be recovered.

Jira:RHEL-114443

Keylime agent no longer fails to enroll with non-RSA certificates

Before this update, the Keylime agent used a single key for both the TLS identity and the payload encryption. As a consequence, when you configured the agent to use a certificate other than RSA, it attempted to use the same key for the payload mechanism and the enrollment process failed.
With this release, the agent relies on two separate keys. As a result, the mutual TLS (mTLS) identity can use alternative cryptographic schemes, and the Keylime agent successfully enrolls with Elliptic Curve Cryptography (ECC) certificates. The payload encryption mechanism still requires a dedicated

RSA key pair.

Jira:RHEL-117122

Keylime agents correctly generate TPM quotes by using ECC keys

Before this update, when generating signed Trusted Platform Module (TPM) quotes, the **keylime-agent-rust** component did not properly support Elliptic Curve Cryptography (ECC) key algorithms. This prevented the agent from generating TPM quote evidence and caused enrollment failures for the ECC key types.
With this update, the **keylime-agent-rust** component correctly handles ECC key algorithms during TPM quote generation. As a result, agents can successfully generate TPM quotes and enroll with the verifier to provide full attestation functionality with ECC keys generated by the TPM.

Jira:RHEL-117441

Keylime verifier correctly validates TPM quotes signed with ECC keys

Before this update, when verifying signed Trusted Platform Module (TPM) quotes from agents, the Keylime verifier component did not properly support Elliptic Curve Cryptography (ECC) key algorithms. This caused attestation failures when agents used the ECC key types **ecc521**, **ecc384**, **ecc256**, **ecc224**, or **ecc192**.
With this update, the verifier correctly handles and verifies TPM quotes signed with ECC keys. As a result, Keylime provides full attestation functionality for these algorithms.

Jira:RHEL-117442

The scp utility correctly handles relative paths containing..

Before this update, the **scp** utility did not expand the **..** parent directory indicator in a path to the directory name. Consequently, **scp** incorrectly handled relative paths containing **...** This update adds special handling for parent directory indicators. As a result, **scp** now processes paths containing **..** correctly.
Jira:RHEL-118406

keylime-policy no longer fails to process remote RPM repositories

Before this update, the **keylime-policy** command failed to close file handles during the analysis of remote RPM repositories, which caused file descriptor leaks. As a consequence, when you used the **--remote-rpm-repo** option to generate a runtime policy, **keylime-policy** failed with a **Too many open files** error. With this update, the command properly closes file handles for all repository metadata and package files and does not exceed the system file descriptor limit.
As a result, **keylime-policy** successfully generates runtime policies from remote RPM repositories.

Jira:RHEL-119028[1]

Restored certificate bundles in /etc/pki/tls and /etc/ssl

Before this update, certificate bundles were removed from **/etc/pki/tls** and **/etc/ssl** as part of the transition to the directory-hash format. Consequently, applications relying on these bundles failed to establish secure connections.
With this update, Red Hat restored the certificate bundles and moved the directory-hash format to RHEL-11. Affected applications can now establish secure connections as before. For RHEL-11 transition guide, see Dropping of cert.pem file.

Jira:RHEL-120696[1]

The **keylime-policy** command correctly handles the **--ima-measurement-list** option

Before this update, if you did not specify a file path for the **--ima-measurement-list** option, the **keylime-policy** command did not properly set the default value. This error blocked other options, such as **--keyrings**, and **keylime-policy** failed to create the runtime policy.
With this update, the **keylime-policy** command uses the default path, **/sys/kernel/security/ima/ascii_runtime_measurements**, when you do not provide a specific value for the **--ima-measurement-list** option. As a result, **keylime-policy** successfully creates the runtime policy.

Jira:RHEL-130158

rust-rpm-sequoia correctly requires OpenSSL 3.5 as a dependency

The **rust-rpm-sequoia** package requires the **openssl** packages in version 3.5, but this was not reflected in the RPM dependency chain. Consequently, you were able to install **rust-rpm-sequoia** without OpenSSL 3.5, but the RPM package management tool subsequently stopped working. With this update, the explicit dependency on OpenSSL 3.5 has been added. As a result, you cannot install **rust-rpm-sequoia** without the required OpenSSL version, which prevents the RPM tool from failing.
Jira:RHEL-130960

/usr/share/*/bin/* binaries work with **fapolicyd**

Before this update, the **fapolicyd** service did not add binaries from **/usr/share/*/bin/** directories to the trust database. For example, the **/usr/share/Modules/bin/mkroot** binary was not added. Consequently, users could not run these binaries when using the **trust=1** option in **fapolicyd** rules. With this fix, the **fapolicyd-filter.conf** file contains ***/bin/***. As a result, you can run binaries from **/usr/share/*/bin/** with the **fapolicyd** service active.
Jira:RHEL-131723

Clevis handles migrations to image mode correctly

Before this update, user and group membership updates from package installations were not properly applied when migrating from package mode to image mode. Consequently, the **clevis** user was not added to the **tss** security group, preventing Clevis from accessing a trusted platform module (TPM) device and retrieving encryption keys during system boot. With this update, the Clevis package installation process is updated to ensure that the **clevis** user is properly added to the **tss** group during image mode updates, even when existing configuration files are preserved. As a result, Clevis can properly access the TPM device and successfully retrieve an encryption key on systems in image mode.
Jira:RHEL-132188

clevis-pin-tpm2 no longer silently ignores invalid JSON

Before this update, the **clevis-pin-tpm2** command did not validate JSON field names during encryption with TPM2 and silently ignored typos and invalid fields, for example, **pcrs_ids** instead of **pcr_ids**. Consequently, users could inadvertently create LUKS bindings with incorrect TPM2 configurations due to typos. This could lead to unlock failures when TPM state changes, potentially making systems unbootable.

This update adds JSON schema validation to reject unknown fields in the TPM2 configuration during encryption. As a result, invalid field names in TPM2 JSON configuration are properly rejected with clear error messages to prevent silent misconfigurations that could cause unlock failures.

Jira:RHEL-138591[1]

SELinux policy update fixes hostname configuration failures

Before this update, a missing SELinux policy rule prevented the **systemd-hostnamed** service from creating a Varlink socket file in the /**run** directory. This issue caused hostname configuration to fail during PXE installations that used Kickstart with **bootc**, which resulted in failed installations.
With this update, the SELinux policy permits the **systemd_hostnamed_t** domain to create the required socket file. As a result, hostname configuration completes successfully.

Jira:RHEL-139385[1]

rust-rpm-sequoia no longer causes RPM to fail for disallowed algorithms

Before this update, when handling signatures with algorithms disallowed by the system-wide cryptographic policies, the **rust-rpm-sequoia** library reported a generic failure error to the RPM package management tool. Consequently, RPM failed to validate signatures on RPM packages with such algorithms. In this update, when **rust-rpm-sequoia** encounters an algorithm disallowed by **crypto-policies**, it reports the **NOTTRUSTED** error message. As a result, you can use **crypto-policies** to disallow one of the algorithms used for signing packages without causing RPM to fail the whole package verification.
Jira:RHEL-144414

12.3. SOFTWARE MANAGEMENT

Review issues that have been fixed for software management in Red Hat Enterprise Linux 10.2.

dnf-automatic can send emails to multiple recipients with default **/usr/bin/mail**

Before this update, if the **dnf-automatic** utility used the **command_email** emitter to send emails to multiple recipients and also used the **/usr/bin/mail** utility installed with the **s-nail** package, **/usr/bin/mail** failed to send an email. With this update, the **dnf-automatic** utility expands the **email_to** keyword in the **command_format** formatting string from a single argument to multiple arguments. As a result, **dnf-automatic** sends emails to multiple recipients with the default **/usr/bin/mail** utility.
Jira:RHEL-94331

RPM no longer fails to install or verify a package with multiple signatures when the package has some **NOTTRUSTED** signatures

Before this update, when you installed or verified a package with multiple signatures, RPM did not correctly determine the overall verification result when the **rpmkeys(8)** utility reported some of the package signatures as **NOTTRUSTED**. A signature can become **NOTTRUSTED** if, for example, its certificate is expired or revoked, or if its algorithm is disabled by system-wide cryptographic policies. As a consequence, RPM failed to install or verify the package even if the package had at least one valid and trusted signature.
This update fixes the verification logic in RPM to correctly handle packages with **NOTTRUSTED** signatures. This update also improves error reporting around this functionality.

As a result, RPM ignores **NOTTRUSTED** package signatures and successfully installs or verifies a package with multiple signatures if the package has at least one valid signature and no invalid signatures. Error messages are also clearer and more accurate when verification actually fails.

Jira:RHEL-112394

DNF no longer fails to install packages that use both supported and unsupported signing algorithms

Before this update, you could not install packages with signatures that used both supported and unsupported package signing algorithms. As a consequence, DNF rejected such packages when verifying their signatures because of the unsupported algorithms. With this update, DNF ignores signatures classified as **NOTTRUSTED** in the **rpmkeys** command output. As a result, DNF can install packages that use both supported and unsupported signing algorithms.
Jira:RHEL-112730

RPM resolves non-local users and groups correctly when installing or verifying packages

Before this update, you could use centralized identity management, for example, through the Lightweight Directory Access Protocol (LDAP), and build a custom package that contains files to be owned by corresponding users or groups. As a consequence, when you installed this custom package, RPM failed to resolve any non-local user and group names, and defaulted to **root** in both cases. This caused files owned by non-local users or groups to be owned by **root** when installed on disk.
With this update, RPM consults the Name Service Switch (NSS) file when resolving user and group names. As a result, as long as the NSS configuration on the system is correct, RPM resolves such non-local users and groups correctly, and the files are owned by the correct User Identifier (UID) and Group Identifier (GID) when installed on disk.

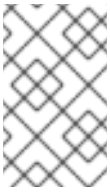

NOTE

If you do not want to use NSS, you can configure the **%_passwd_path** or **%_group_path** macro. RPM will only use the specified local **passwd(5)** and **group(5)** files when resolving names.

NOTE

When installing or verifying packages in an alternate root directory specified with the **--root** option, RPM only uses the **passwd(5)** and **group(5)** files, or the paths configured with the **%_passwd_path** and **%_group_path** macros, inside the target **root** directory and does not consult NSS at all.

For more information, see the **--root** description in the **rpm(8)** man page.

Jira:RHEL-118365

DNF correctly performs comparison of**epoch-version-release** for upgrade transactions

Before this update, DNF incorrectly performed comparison of the **epoch-version-release** (**EVR**) RPM package information. As a consequence, if you performed two subsequent upgrade transactions for a package that had the same **epoch-version** but different **release**, DNF identified the overall transaction as a downgrade. This update fixes the **EVR** comparison. As a result, DNF identifies two subsequent package upgrades with different release versions as an upgrade.

Jira:RHEL-128443[1]

DNF no longer attempts to automatically remove protected packages installed as dependencies

Before this update, if you installed a protected package as a dependency required by only one other package and had the **clean_requirements_on_remove** configuration option enabled, DNF failed to perform any transaction that tried to remove the protected package if this package became an unused dependency. This prevented the removal of the package that depended on it, because DNF would automatically attempt to remove the protected dependency as well. With this update, DNF treats all protected packages as explicitly installed by the user. As a result, DNF no longer attempts to automatically remove protected packages, allowing the removal of the package that depends on it.

Jira:RHEL-128445[1]

12.4. SHELLS AND COMMAND-LINE TOOLS

Review issues that have been fixed for shells and command-line tools in Red Hat Enterprise Linux 10.2.

Fixed the ipmievd service start failure due to missing PID file

Before this update, the **ipmievd** systemd service failed to start because the service did not create the required PID file during initialization. As a result, the service could not store its process ID and timed out. With this update, the service configuration creates the PID file before starting **ipmievd** to ensure that the service starts correctly.

Jira:RHEL-112449[1]

volume_key successfully retrieves backup passphrases in FIPS mode

Before this update, the **volume_key** utility used functions that were incompatible with Federal Information Processing Standards (FIPS) when retrieving a backup passphrase from an escrow packet. Consequently, **volume_key** failed and reported an error on systems with FIPS mode enabled. This update ensures that the backup passphrase retrieval function is FIPS-compliant. As a result, you can successfully retrieve backup passphrases on FIPS-enabled systems.

Jira:RHEL-146218[1]

12.5. NETWORKING

Review issues that have been fixed for networking in Red Hat Enterprise Linux 10.2.

The NetworkManager sriov.vfs property supports thereapply operation

Before this update, NetworkManager could not dynamically apply changes if a user changed the **sriov.vfs** property. As a consequence, NetworkManager connections with Single Root I/O Virtualization (SR-IOV) settings required a restart after modifications. With this release, **sriov.vfs** now supports the **reapply** operation if the total number of virtual functions (VFs) does not change. As a result, restarting a connection after modifying SR-IOV settings is no longer required in the mentioned scenario.

Jira:RHEL-95844

The xdp-trafficgen utility works correctly on ARM systems

Before this update, the **xdp-trafficgen** utility failed on ARM systems with a **Missing required option '--interface'** error even if you specified the **-i <interface>** option. As a consequence, it was not possible to probe eXpress Data Path (XDP) support on a specific interface. This update fixes the problem, and the **-i <interface>** option works correctly on ARM systems.

Jira:RHEL-105793

NetworkManager clients can set a global-level DNS search domain without defining a DNS server

Before this update, if a client, such as the Nmstate API or the GNOME control center application, used the D-Bus API for changes on a global level, it was not possible to set DNS search domains without defining a DNS server. This update fixes the problem, and clients can define only a global-level DNS search domain.
Jira:RHEL-109853

NetworkManager-wait-online.service is no longer a hard requirement in **nmstate.service**

Before this update, the **nmstate.service** systemd unit had a strict **Requires** dependency on the **NetworkManager-wait-online.service** unit. Consequently, in environments where **NetworkManager-wait-online** failed or timed out, the **nmstate** state service failed to start and Nmstate could not apply the necessary network configurations. This update replaces **Requires** in the unit with **Wants**. As a result, the **nmstate** service starts regardless of the status of **NetworkManager-wait-online**, and Nmstate can apply network configurations.
Jira:RHEL-114959

Unbound follows system-wide crypto-policies for TLS

Before this update, the Unbound package explicitly disabled TLS 1.2 on server sockets for DNS over TLS (DoT). Consequently, servers could not use TLS 1.2 regardless of system settings.
With this update, the default system-wide crypto-policies manage protocol disabling . As a result, TLS server sockets accept older protocols like TLS 1.2, and TLS 1.1 in **LEGACY** policy mode.

Jira:RHEL-147790

12.6. KERNEL

Review issues that have been fixed for kernel in Red Hat Enterprise Linux 10.2.

The USB bulk transport path no longer leaks USB protocol bytes

Before this update, a memory leak of USB protocol data in the USB bulk transport path could expose bytes from the USB protocol into user space when devices incorrectly skipped the data phase. This behavior was detected by the Linux Test Project (LTP) **ioctl_sg01** test and indicated that the SCSI request block (SRB) transfer buffer was not cleared in these situations.
With this update, the RHEL kernel is updated to clear the SRB transfer buffer when the data phase is incorrectly skipped. As a result, the USB bulk transport path no longer leaks USB protocol bytes into the user space, and the **ioctl_sg01** test now completes successfully.

Jira:RHEL-2588[1]

Incorrect state decoding in **perf_sched** fixed, and the **perf_sched** test suite now passes

Before this update, the incorrect state decoding in **perf_sched** caused the perf tool test suite to fail. This affected the perf tool test suite results. With this release, a patch for correct process state decoding in **perf_sched** test has been implemented. As a result, the **perf_sched** test suite now passes.
Jira:RHEL-68347[1]

Update resolves ftrace testing failures for uprobe events by using readelf for entry point determination

Before this update, the **uprobe** events test during **ftrace** testing failed due to an issue with entry point determination. This fix involves using readelf for entry point determination in the **uprobe** tests. As a result, **ftrace** testing failures for **uprobe** events and Kprobe event parsing are resolved, improving **ftrace** test stability.

Jira:RHEL-87219[1]

12.7. FILE SYSTEMS AND STORAGE

Review issues that have been fixed for file systems and storage in Red Hat Enterprise Linux 10.2.

SCSI tape driver fix now enables device ID IOCTLs after reset

Before this update, a regression in the SCSI tape (**st**) driver caused certain tape applications to fail after a device reset (such as a third-party power-on reset). When these applications attempted to verify device ID information by using **ioctl** commands, the driver blocked the request if the buffer state was not ready. This caused errors such as "device /dev/nst1 failed on scsi ioctl(idlun)" and the affected tapes entered an error state.
With this update, the **st** driver has been fixed to ensure that informational **ioctl** commands, such as **idlun**, can execute regardless of the internal buffer state. As a result, tape applications can now successfully verify device information after a reset.

Jira:RHEL-115965[1]

Multipath persistent reservation handling is now more robust and consistent

Before this update, the **libmpathpersist** library, which is used by the **mpathpersist** command, had several issues and corner cases that affected persistent reservation handling for multipath devices. This caused the following problems:

- Numerous **mpathpersist** operations failed on a multipath device.
- Persistent reservations sometimes ended up in an inconsistent state. As a consequence, the multipath device denied write access when it was supposed to be allowed, and allowed write access when it was supposed to be prohibited.

With this release, multiple areas of **libmpathpersist** have been redesigned and fixed to ensure correct and consistent behavior. As a result, **mpathpersist** commands on multipath devices now work the same as the equivalent **sg_persist** commands on SCSI devices. I/O access to multipath devices also consistently reflects the device's persistent reservation state.

Jira:RHEL-118720[1]

The Anaconda installer can now use iSCSI LUNs with ID 256 or higher

Before this update, starting an operating system installation on a system that used iSCSI storage could cause the Anaconda installer to crash. This occurred when the iSCSI Logical Unit Number (LUN) ID was 256 or higher.
This update includes a fix to the LUN ID parsing logic in the **blivet** library. As a result, installations on systems that use iSCSI targets with LUN IDs of 256 or greater can now proceed.

Jira:RHEL-122305

Added a new VDO **vdocalculatesize** utility and improved metadata error handling

This update introduces the **vdocalculatesize** utility. The **vdocalculatesize** computes Virtual Data Optimizer (VDO) volume size and memory requirements based on parameters such as logical size, physical size, slab size, index memory size, and block map cache size. As a result, you can accurately plan and provision VDO volumes, reducing configuration uncertainty for VDO storage deployment. This release also fixes error handling for scenarios in which VDO metadata becomes corrupted.

Jira:RHEL-129906

multipathd logs offline path warnings for uninitialized paths

Before this update, if **multipathd** started or reconfigured while a path was offline, the daemon did not print regular offline warnings for that path. This made it difficult to identify issues with uninitialized paths.
With this update, **multipathd** prints offline messages for uninitialized paths. As a result, you can monitor path status consistently.

Jira:RHEL-133815[1]

Fixed NVMe **subsystem reset** recovery on PowerPC

Before this update, issuing the **nvme subsystem-reset** command on the PowerPC platform caused the Non-volatile Memory Express (NVMe) device to enter the **resetting** state and it failed to recover. As a consequence, the device hung and required a system reboot to recover.
With this release, the NVMe device recovers correctly after a **subsystem reset**. It is temporarily inaccessible while transitioning from the **resetting** state to the **live** state.

Jira:RHEL-137767[1]

12.8. HIGH AVAILABILITY AND CLUSTERS

Review issues that have been fixed for high availability and clusters in Red Hat Enterprise Linux 10.2.

Nodes no longer unexpectedly leave the cluster after rejoining

Before this update, when a node left a cluster, the cleanup of its transient attributes was handled by two separate components. As a consequence, a node's shutdown attribute might not have been cleared before the node attempted to rejoin the cluster, causing the node to immediately leave again.
With this release, the responsibility for clearing all transient node attributes has been consolidated into a single component.

As a result, these timing issues are no longer possible, and nodes can rejoin the cluster without being immediately removed due to stale **shutdown** attributes.

Jira:RHEL-23082

Warning messages added when using the **-f** flag to modify CIB files

Before this update, when a user executed **pcs resource delete**, **pcs cluster node remove-remote**, or **pcs booth remove** using the **-f** flag to modify a CIB file directly, **pcs** would perform the deletion but silently omit cleanup actions that require a live cluster, such as stopping resources or removing nodes from Pacemaker.

With this update, warning messages are displayed whenever live cluster cleanup actions are skipped due to the use of the **-f** flag.

As a result, users are alerted that they must perform manual cleanup actions on the live cluster when modifying configuration files offline. Note that the usage of the **--force** flag to skip resource stopping in these commands has been deprecated and will be removed in a future release. The **--force** flag retains its functionality to override validation errors. Users should now use the **--no-stop** flag to explicitly skip resource stopping before deletion.

Jira:RHEL-76157

The pcs resource delete command blocks the deletion of running unmanaged resources

Before this update, if you attempted to delete an unmanaged resource by using **pcs resource delete** while that resource was still running, the resource would be removed from the Cluster Information Base (CIB) but remain active in the running configuration. This left the resource in an **ORPHANED** state, which could lead to cluster instability and resource management issues.
With this update, **pcs** returns an error if a deletion request includes any unmanaged resources that are currently running.

As a result, **pcs resource delete** prevents the creation of orphaned resources by requiring that they be stopped before deletion.

Jira:RHEL-76162

Resource and stonith agent descriptions retain original formatting

Before this update, **pcs** automatically wrapped resource and stonith agent descriptions to fit within the terminal window. Consequently, any formatting done by the agents' authors-such as new lines, paragraphs, lists, or tables-was removed, often making the descriptions difficult to read.
With this update, **pcs** no longer reformats the description text.

As a result, **pcs** displays resource and stonith agent descriptions exactly as the agents' authors intended, preserving the original structure and improving readability.

Jira:RHEL-111451

The db2 resource agent handles reintegration correctly

Before this update, the **db2** resource agent could encounter a race condition when a node was reintegrating into the cluster. Consequently, the reintegrating node could incorrectly attempt to start as a "Primary" instance.
With this update, a "reintegration" attribute has been added to the agent. This allows the agent to correctly identify whether it is expected to join as a "Primary" or not, avoiding the race condition.

As a result, reintegration works correctly. Note that in order to prevent issues during the upgrade, you must disable all **db2** resources before applying the update and re-enable them only after the update is complete on all nodes.

Jira:RHEL-115495[1]

12.9. COMPILERS AND DEVELOPMENT TOOLS

Review issues that have been fixed for compilers and development tools in Red Hat Enterprise Linux 10.2.

Fix missing gzip dependency for compressed locale character maps in glibc-locale-source

Before this update, the **glibc-locale-source** package provided character maps in **gzip** compressed format but did not declare a dependency on the **gzip** package. As a consequence, using **localedef** with a character map provided by **glibc-locale-source** could fail if **gzip** was not installed on the system because the compressed archive could not be uncompressed.
With this release, **glibc-locale-source** now depends on the **gzip** package. This change ensures the required compression utility is present, allowing **localedef** to process character maps correctly. As a result, using **localedef** with character maps provided by **glibc-locale-source** now works as expected even on systems where **gzip** was previously missing.

Jira:RHEL-102553

The glibc exit function no longer crashes on simultaneous calls

Before this update, simultaneous calls to the glibc **exit** function and concurrent **stdio.h** stream operations in multi-threaded applications were not synchronized. As a consequence, applications could terminate unexpectedly or experience data corruption. With this update, the **exit** and **quick_exit** functions synchronize **stdio.h** stream flushing and allow only one exit call to proceed. As a result, applications no longer crash in this scenario, and overall reliability is improved.
Applications that perform blocking read operations on **stdio.h** streams, such as with **getchar**, or that use **flockfile** to lock streams, cannot exit until the read operation completes or the lock is released. This behavior is required by the POSIX standard.

Jira:RHEL-111117[1]

glibc now returns complete group membership results when NSS group merges fail with ERANGE

Before this update, when looking up group membership on systems where Name Service Switch (NSS) merged groups from more than two services, a merge between two groups that failed due to an insufficient internal buffer caused glibc to skip the merge result instead of retrying the operation with a larger buffer.
As a consequence, on systems with more than two group database sources, querying group information, such as with the **getent** group command, produced incomplete or empty group membership results in some cases.

With this update, glibc no longer skips merge failures that are caused by an insufficient internal buffer and instead retries the merge with a larger buffer as intended.

As a result, group membership lookups on systems with multiple group database sources now return complete and correct group membership data.

Jira:RHEL-114265[1]

Fixed uninitialized result from sem_open when opening missing semaphore

Before this update, calling the **sem_open** function for a named semaphore that did not exist and without specifying the **O_CREAT** flag could return an uninitialized value instead of a defined error indicator.
As a consequence, affected applications observed undefined behavior, such as attempting to use an invalid semaphore handle and misinterpreting the failure because errno was not set to a meaningful value.

With this release, **sem_open** explicitly returns **SEM_FAILED** and sets **errno** to **ENOENT** when it is called for a semaphore that does not exist and the **O_CREAT** flag is not specified.

As a result, applications reliably detect this error condition and can handle missing semaphores in a predictable and standards-compliant way.

Jira:RHEL-119392[1]

glibc stdio flushing issues fixed for input streams and shared file descriptors

Before this update, the glibc standard I/O implementation did not fully comply with POSIX when flushing input streams. This caused **fflush** to mishandle input streams after **ungetc**, inconsistent behavior when called as **fflush(NULL)**, and incorrect file offsets when **fclose** operated on shared file descriptors or special character devices.
As a consequence, applications might observe unexpected input stream state, incorrect underlying file positions, and file positioning errors when using **fseek** and **fflush** on memory-mapped input files, which can lead to misreads or subtle data-processing bugs.

With this release, the **glibc** stdio library is corrected so that **fflush** handles input streams in a POSIX-compliant way, including after **ungetc** and when invoked as **fflush(NULL)**. In addition, **fclose** now updates the underlying file offset for shared file descriptors and works correctly with special character devices, and the file positioning logic for **fseek** and **fflush** on memory-mapped input files is fixed.

As a result, applications that rely on **stdio** for input processing, shared file descriptor usage, or memory-mapped input files now behave predictably and correctly after the update.

Jira:RHEL-119434[1]

glibc NSS database lookup stability improvement

Before this update, missing checks in the **__nss_database_get** function in the **glibc** package could cause null pointer dereferences and assertion failures during Name Service Switch (NSS) database lookups. As a consequence, applications relying on NSS could terminate unexpectedly, or the C library could crash under specific lookup conditions.
With this release, additional validation checks are added to the NSS database lookup path in **glibc** to handle invalid or unexpected internal states safely. As a result, NSS database lookups are more robust, and system stability is improved.

Jira:RHEL-150270

Duplicate DNS queries fixed when the search path is set to

Before this update, when the Domain Name System (DNS) search path in **/etc/resolv.conf** file contained a single **.** entry, the **glibc** DNS stub resolver queried both the original domain name and the same domain name with a trailing dot.
As a consequence, DNS queries for non-existent domains were duplicated, increasing the load on DNS servers.

After this update, the **glibc** DNS stub resolver no longer appends a trailing dot to domain names when the search path contains only a single **.** entry.

As a result, DNS queries are no longer duplicated in this configuration, reducing unnecessary DNS traffic and server load.

Jira:RHEL-142675

12.10. IDENTITY MANAGEMENT

Review issues that have been fixed for Identity Management (IdM) in Red Hat Enterprise Linux 10.2.

Directory Server tools consistently accept unit suffixes when configuring the LMDB database maximum size

Before this update, **dscreate** and **dsconf** used different functions to parse and display the LMDB database maximum size (**nsslapd-mdb-max-size**). As a consequence, **dscreate create-template** displayed the value as a raw floating-point number in bytes, while **dsconf backend config set --mdb-max-size** accepted values in bytes only, making it difficult to configure consistent values across the two tools.
With this update, both tools use the same parsing functions and accept values with unit suffixes (**k**, **m**, **g**, **t**), automatically aligning the result to the nearest page boundary. As a result, administrators can use human-readable size values consistently across **dscreate** and **dsconf** when setting the LMDB database maximum size.

Jira:RHEL-64019

The Directory Server web console displays sub-suffixes whose parent suffix is a regular entry

Before this update, the Directory Server web console only displayed sub-suffixes whose **nsslapd-parent-suffix** attribute exactly matched an existing backend suffix. As a consequence, sub-suffixes with a parent suffix pointing to a regular LDAP entry (rather than a backend suffix) were not visible in the console's suffix tree, even though they appeared correctly in the **dsconf backend suffix list** output.
With this update, the web console correctly identifies sub-suffixes that fall under a backend suffix, regardless of whether the parent suffix is a backend suffix itself. As a result, all configured sub-suffixes are displayed in the web console suffix tree.

Jira:RHEL-76835

Directory Server no longer fails at shutdown when the retro changelog trimming thread is active

Before this update, the retro changelog plugin's internal lock object was freed while the trimming thread was still holding a reference to it when **ns-slapd** started shutting down. As a consequence, the server could fail with a segmentation fault.
With this update, the server waits for all active plugin threads to finish before freeing plugin resources during shutdown. As a result, **ns-slapd** shuts down cleanly even when retro changelog trimming is in progress.

Jira:RHEL-86312

LDAP searches with a single component in compound filters return correct results

Before this update, Directory Server did not correctly evaluate compound LDAP filters that contained only a single filter component, such as **(&(cn:dn:=groups))**. As a consequence, group search queries using these filters returned no results, causing failed group lookups and potentially incorrect access control. With this update, filter evaluation logic is updated to correctly handle compound filters with a single component. As a result, existing group search filters such as **(&(cn:dn:=groups))** return the expected entries, restoring predictable LDAP behavior for applications and scripts.

Jira:RHEL-89601

User resolution no longer fails if **name** ID user overrides exist for IdM AD users

Before this update, when a **name** ID user override existed for IdM AD trusted users, user resolution failed because the auto private group could not be resolved. With this update, the IdM provider retries to fetch the user object if no group override is found. As a result, the auto private group of *<overwritten_name>*@ad.domain can be resolved, and user resolution succeeds.

Jira:RHEL-94545[1]

Directory Server ignores **memberOfDeferredUpdate** setting on instances with LMDB

Before this update, the **memberOfDeferredUpdate** configuration attribute, which is only effective for a Berkeley DB (BDB) backend, was not ignored on instances with a Lightning Memory-Mapped Database Manager (LMDB) backend. As a consequence, if **memberOfDeferredUpdate** was enabled on an LMDB instance, the Directory Server could become unresponsive during MemberOf plugin processing of large or complex groups.
With this update, Directory Server ignores the **memberOfDeferredUpdate** setting on instances with LMDB. As a result, processing large or complex groups no longer causes the server to become unresponsive.

Jira:RHEL-106502

dsctl db2index no longer reindexes all attributes when specific attributes are requested

Before this update, running **dsctl db2index** with the **--attr** option but without specifying a backend name caused the **--attr** option to be silently ignored. As a consequence, all attributes across all backends were reindexed instead of only the specified ones, which could take a significant amount of time on large databases.
With this update, **dsctl db2index** requires a backend name as a positional argument, and the **--attr** option correctly limits reindexing to the specified attributes for the given backend. As a result, only the requested attributes are reindexed when a backend name and the **--attr** option are both provided.

Jira:RHEL-111220[1]

The MemberOf fixup task completion message correctly displays the membership attribute name

Before this update, when the MemberOf plugin completed a global fixup task, the plugin freed its configuration structure before logging the completion message. As a consequence, the completion log message displayed (**null**) instead of the membership attribute name.
With this update, the MemberOf plugin logs the fixup task completion message before freeing its configuration structure, ensuring the attribute name is available when the message is written. As a result, the completion log message displays the correct membership attribute name, making it easier for administrators to verify fixup operations and troubleshoot issues.

Jira:RHEL-117520[1]

The Directory Server web console no longer fails with an error when enabling replication on a consumer

Before this update, when enabling replication on a consumer, the **dsconf** utility printed a warning about changelogs to the **stdout** stream instead of **stderr**. As a consequence, the textual warning broke JSON parsing in the Directory Server web console, which expects pure JSON on **stdout**.
With this update, **dsconf** utility was updated so that the warning about changelogs on consumer replicas is written to **stderr**. As a result, the Directory Server web console successfully loads the **Replication** tab after enabling replication on a consumer or changing a role to consumer.

Jira:RHEL-122674

New **notes=N** and **notes=B** search indicators to identify asynchronous operations in the Directory Server access log

Before this update, asynchronous requests that exceeded the maximum number of threads per connection caused server unresponsiveness without identification in the Directory Server access logs. As a consequence, it was difficult to diagnose server unresponsiveness.
With this release, Directory Server uses the new search indicators in the access logs to identify such requests: **notes=N** defines that the operation is not synchronous. **notes=B** defines that the operation blocks other new incoming operations: pending operations, not the read operations, are delayed.

In both cases, you might need to increase the **nsslapd-maxthreadsperconn** attribute value to allow a connection to use more threads.

Jira:RHEL-123220

Online initialization of a Directory Server consumer no longer fails with a **LDAP_BUSY** error

Before this update, the replication agreement could send entries faster than the consumer was able to import during online initialization. In that situation, the consumer responded with an **LDAP_BUSY** error. As a consequence, the replication agreement did not handle this error and terminated the online initialization.
With this update, the replication agreement handles received **LDAP_BUSY** responses by retrying the operation after a delay. As a result, online initialization completes successfully even when the consumer temporarily cannot keep up with the rate of incoming entries.

Jira:RHEL-123663[1]

LDAP searches with spaces in DN filter values no longer return incorrect results

Before this update, a regression in the handling of filters containing distinguished name (DN) caused LDAP searches with spaces inside DN values in the filter, such as **(member=uid=user, ou=people,dc=example,dc=com)**, to be evaluated incorrectly. As a consequence, applications received incomplete group membership and search results.
With this update, Directory Server normalizes and correctly compares DN values in the filter, accepting filters both with and without spaces in DN components. As a result, LDAP searches that include spaces in DN values return the same, complete results as in earlier RHDS versions, restoring expected application behavior.

Jira:RHEL-123664[1]

Directory Server deletes access logs as expected

Before this update, when access log compression was enabled, the log rotation logic failed to correctly recognize **.gz**-suffixed rotated access log filenames while rebuilding the internal rotation

information, so compressed logs were not associated with their corresponding rotation entries. As a consequence, the **nsslapd-accesslog-list** did not contain the actual files on disk, and access logs accumulated until manual cleanup was required to prevent disks from filling.
With this update, the log rotation logic was updated to correctly parse and match rotated access log filenames regardless of whether they are compressed (with a **.gz** suffix) or uncompressed, ensuring compressed logs are included when rebuilding rotation information and validating previous log files. As a result, compressed rotated access logs are properly tracked and removed according to the configured rotation settings.

Jira:RHEL-124694

Online initialization of large databases progresses as expected

Before this update, when initializing replication with very large databases, especially after major subtree moves, the initialization could appear stalled after sending the initial suffix entry, because it spent excessive time building and checking large internal ID lists. As a consequence, the server experienced long CPU spikes, initialization was delayed or incomplete, and replicas remained outdated for an extended period.
With this update, the internal ID list lookup logic used during online initialization was optimized, making it scalable even with very large datasets. As a result, replication online initialization progresses as expected on large databases.

Jira:RHEL-128906

Replication no longer fails with**Can't locate CSN** errors after an offline import

Before this update, when a replica was reinitialized by using an offline import, the replication keep-alive update was triggered before the replica had time to synchronize with the other suppliers. As a consequence, **Can't locate CSN** (Change Sequence Number) errors were logged and some changes were not replicated to consumers.
With this update, the initial delay before the first keep-alive update matches the value of the **nsds5ReplicaKeepAliveUpdateInterval** attribute, which defaults to 1 hour, and a warning is displayed if this interval is less than the maximum backoff timer. As a result, the replica has sufficient time to synchronize from other suppliers after a reinitialization, and replication proceeds without CSN errors.

Jira:RHEL-129675[1]

Directory Server database initialization no longer fails with an**MDB_BAD_VALSIZE** error

Before this update, when indexing an attribute, Directory Server erroneously extended the prefix of the index key. The more values were indexed, the longer the prefix became. Adding entries with large values accelerated the issue, because the server also appended a hash to the key. For example, entries in a FreeIPA deployment with many certificates triggered an **MDB_BAD_VALSIZE** error. As a consequence, key sizes could exceed the LMDB maximum key size, and Directory Server could not initialize the database during import or replication when the dataset contained such entries.
With this update, Directory Server corrects the index key handling to prevent the **MDB_BAD_VALSIZE** condition. As a result, database initialization succeeds when importing or replicating datasets that contain entries with large numbers of long indexed attribute values.

Jira:RHEL-133085

Directory Server no longer fails under heavy operations involving the NDN cache

Before this update, a defect in the concread dependency used by the Named Data Networking (NDN) cache caused LinCowCell chain drops to incorrectly free shared links when multiple references existed to the same chain. As a consequence, under heavy operations involving the NDN cache, the server could hit a use-after-free condition and fail with a segmentation fault in **atomic_compare_exchange()**, leading to erratic downtime.
With this update, the **389-ds-base** package uses concread version 0.5.10, which correctly stops freeing data when a shared cache link is detected. As a result, NDN cache operations are handled safely, preventing the segmentation fault.

Jira:RHEL-138729

Resolved DNS record creation failure when reverse zone is missing

Before this update, the **ipadnsrecord** module in **ansible-freeipa** ignored the **create_reverse** parameter. As a consequence, when users attempted to add **A** or **AAAA** records, the module incorrectly always required an existing reverse DNS zone and the task failed with a "DNS zone not found" error.
With this release, the module logic verifies the status of the **create_reverse** flag before attempting to validate or locate a reverse zone and skips the check entirely if it is set to **false**. As a result, the **ipadnsrecord** module successfully adds **A** and **AAAA** records to IdM-managed zones without requiring an existing reverse zone when **create_reverse** is set to **false**.

Jira:RHEL-140606

12.11. SSSD

Review issues that have been fixed for SSSD in Red Hat Enterprise Linux 10.2.

adcli correctly identifies machine account principals in multi-realm keytabs

Before this update, when connecting to a domain to update a password, **adcli** always used the Kerberos realm of the first entry in the keytab file. As a consequence, on systems where the keytab contained multiple realms, the renewal process failed with a "no suitable keys" error if the required realm was not listed first. With this release, **adcli** searches the keytab for a principal that matches the target domain. As a result, machine account password renewals now succeed regardless of the order of entries in the keytab.
Jira:RHEL-2518

adcli testjoin correctly identifies the joined domain in multi-principal keytabs

Before this update, the **adcli testjoin** command unconditionally used the domain or realm from the first entry found in the keytab file to perform its diagnostic test. As a consequence, on systems where the keytab contained principals from multiple domains, **adcli testjoin** would often attempt to connect to an incorrect domain and fail with a "Realm not local to KDC" error.
With this release, **adcli** uses the realm from the keytab as the domain name when the domain is not explicitly specified. As a result, users can reliably verify domain connectivity without encountering false authentication failures.

Jira:RHEL-5044

User creation fails with invalid**sAMAccountName** input

Before this update, user creation with, for example, a User Principal Name (UPN) format that

includes the @ character instead of a **sAMAccountName** attribute, caused **adcli** to create user objects with a **sAMAccountName** which contained invalid characters. As a consequence, Active Directory (AD) operations involving that user could break. With this release, **adcli** validates the input string for user creation against a list of illegal characters before attempting to create the entry. As a result, **adcli** terminates user creation if the input is not a valid **sAMAccountName** value. This prevents the creation of malformed user objects and ensures smoother AD operation.

Jira:RHEL-5050[1]

12.12. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Review issues that have been fixed for Red Hat Enterprise Linux system roles in Red Hat Enterprise Linux 10.2.

The **network** RHEL system role no longer fails to look up routing tables by name

The **/usr/share/iproute2/rt_tables** file contains certain built-in routing table names, such as **main**. Before this update, if an administrator used the **network** RHEL system role to modify the routing table and specified a routing table by its name in a playbook, the role failed with the following error:

```
cannot find route table main in /etc/iproute2/rt_tables or /etc/iproute2/rt_tables.d/
```

With this update, the **network** RHEL system role no longer fails to look up routing tables by name in **/etc/iproute2/rt_tables** and files in the **/etc/iproute2/rt_tables.d/** directory.

Jira:RHEL-110865[1]

Storage role no longer fails when **/etc/fstab** is missing

Before this update, the storage role crashed on systems where **/etc/fstab** was absent. As a consequence, systems without a file system table configuration experienced failures.
With this update, the storage role checks whether **/etc/fstab** exists before attempting to parse it. As a result, systems without this file no longer experience a crash when using the storage role.

Jira:RHEL-115033

External configuration files correctly override all the **sshd_config** options

Before this update, external configuration files were not loaded first, which prevented overrides of all options in the **sshd_config** file. Consequently, users experienced incorrect OpenSSH daemon configuration. With this update, external configuration files take priority. As a result, users can override all options in the **sshd_config** file.
Jira:RHEL-123016

The **network** RHEL system role no longer reports an incorrect state when removing profiles

Before this release, when you used the **network** RHEL system role with the **persistent_state: absent** setting to remove undefined profiles, the role attempted to delete the loopback interface profile. Because the system automatically recreates this profile immediately, Ansible incorrectly reported a **changed** state. This bug fix adds the loopback device to the role-internal **black_list_names** variable. As a result, the **network** RHEL system role ignores the loopback interface. This prevents unnecessary changes and the role reports an **ok** state.
Jira:RHEL-123026

Fixed ZeroDivisionError when creating LVM volumes without a specified size

Before this update, creating an LVM volume without specifying a size could cause a ZeroDivisionError. This occurred because the **blivet** module treated a volume with no specified size as zero.
With this release, if you do not specify size, the volume uses all available space in the pool. As a result, LVM volumes are created successfully even when a size is omitted.

Jira:RHEL-123523

The **nbde_client** role correctly maintains idempotence after failed binding operations

Before this update, when the **nbde_client** system role failed to add a required binding to a LUKS-encrypted volume, the rollback mechanism did not always function correctly. This caused idempotence issues, where subsequent attempts to run the role would fail or produce unexpected results because the system was left in a partially modified state.
With this update, the role performs a backup of the LUKS header before initiating any binding operations. If an operation fails, the role uses this backup to restore the header to its original state. As a result, the role correctly maintains idempotence and ensures the system remains in a consistent state even if a binding fails to be added.

Jira:RHEL-128428[1]

The **aide** system role supports dynamic database configuration for multiple AIDE versions

Before this update, the **aide** system role used the deprecated **database** variable in its templates. On systems running Advanced Intrusion Detection Environment (AIDE) version 0.17 or later, including RHEL 10.2, RHEL 9.8, and CentOS Stream 9, this caused the AIDE service to fail during configuration parsing.
With this update, the role introduces the **database_in** and **aide_version** variables to dynamically detect the installed AIDE version and apply the appropriate configuration syntax automatically.

As a result, the **aide** system role provides consistent file integrity monitoring across different releases without requiring manual configuration changes.

Jira:RHEL-129309

Improved error handling for empty disk lists in **blivet**

Before this update, the code failed to check if the disks list was empty before accessing **disks[0]** in the **blivet** module. As a consequence, an unhandled **IndexError** caused playbook failures, leading to poor performance.
With this update, the module checks whether the disk list is empty before accessing it. If no disks are available, a clear error message is displayed instead of triggering an exception.

Jira:RHEL-137261

vpn role generates valid **ipsec.conf** file for unmanaged hosts

Before this update, when you tried to generate an **ipsec.conf** file for VPN connection between managed and unmanaged hosts, a logic error in the Ansible Playbook caused the task to fail. With this update, the Ansible Playbook references the host and subnet information correctly.
As a result, the **vpn** system role generates a valid **ipsec.conf** file for this scenario.

Jira:RHEL-145219

The selinux system role supports static imports even when some variables are undefined

Before this update, undefined variables, such as module paths, caused the **selinux** system role to fail during template expansion if the **import_role** directive was used. This occurred because Ansible attempts to resolve variables in task **name** fields immediately, even if those tasks are within a block with a **when** condition that evaluates to false.
With this update, task names use the **default**, or **d**, filter to provide a fallback value for potentially undefined variables. This ensures that static imports succeed without error, and dynamic usage with the **include_role** module still provides detailed task information when variables are present.

As a result, the **selinux** role functions correctly in playbooks that use the **import_role** directive even when no specific module path is defined.

Jira:RHEL-145247

The firewall RHEL system role installs NetworkManager on managed nodes in order for PCI interface ID lookups to work correctly

Previously, if you wanted to look up the interface name by specifying the PCI id for the interface by using the **interface_pci_id** parameter, and NetworkManager was not installed, the **firewall** RHEL system role was unable to look up the interface by PCI ID and displayed a warning. As a consequence, the role failed to configure the **firewalld** service by using the specified **interface_pci_id** variable. With this update, the role ensures that NetworkManager is installed, and the **firewall** RHEL system role works as expected.
Jira:RHEL-150780

Resolved task name expansion issues in Ansible roles

Before this update, if you used **import_role** with modules that had no path set, the role issued undefined variable errors. This occurred because Ansible attempted to expand templates in task names within a **block** regardless of the **when** conditions.
With this update, the **d** filter provides a default value for these variables. As a result, the role no longer errors with **import_role** and modules without a defined path, and continues to provide additional context in task names when used with **include_role**.

Jira:RHEL-150788

Loop mount errors on RHEL 7 are resolved

Before this update, the **blivet** module called an undefined function during loop mounts on Red Hat Enterprise Linux 7 because the **libblockdev-loop** package was missing. As a consequence, the role failed with the "The function 'bd_loop_get_backing_file' called, but not implemented" error.
With this update, the **libblockdev-loop** package is installed, which prevents **blivet** errors during loop mounts on RHEL 7.

Jira:RHEL-151437

12.13. VIRTUALIZATION

Review issues that have been fixed for virtualization in Red Hat Enterprise Linux 10.2.

Post-copy migration no longer causes connection issues on IBM Z

After migrating a virtual machine (VM) between IBM Z hosts by using post-copy migration, the VM

previously in some cases lost network connection and required resetting its network interface to reconnect. With this update, the kernel handles post-copy initiation properly, and the problem no longer occurs.
Jira:RHEL-42486

virtiofsd no longer exhausts open file descriptors when sharing directories with many files

Before this update, **virtiofsd** used file descriptors to hold references to files in a virtiofs-shared directory until the guest kernel invalidated its cache. As a consequence, when accessing a large number of files through **virtiofs**, **virtiofsd** accumulated open file descriptors and exceeded the system limit. This caused commands such as **rsync** and **du** to fail with **Too many open files** errors and in some cases caused **virtiofsd** to crash.
With this update, **virtiofsd** defaults to using inode file handles instead of file descriptors to hold references to files. As a result, **virtiofsd** no longer exhausts the open file descriptor limit when working with **virtiofs**-shared directories that contain a large number of files.

Jira:RHEL-99895[1]

Live migration of VMs with multiple CPU threads no longer fails with a CPU feature mismatch

Before this update, the **libvirt** package reported the **ht** (Hyper-Threading) CPU feature flag inconsistently between the source and destination hosts during live migration. As a consequence, live migration of virtual machines (VMs) that were configured with multiple CPU threads could fail with the following error:

```
guest CPU doesn't match specification: extra features: ht
```

With this update, the **libvirt** package correctly handles the **ht** CPU feature flag during migration. As a result, VMs configured with multiple CPU threads can be successfully migrated between hosts.

Jira:RHEL-104216

TDX attestation no longer requires rebooting the host

Previously, after you installed the **linux-sgx** packages on your host, Intel Trust Domain Extensions (TDX) attestation on your virtual machines (VMs) only worked after you rebooted the host. Now, the **/dev/sgx_provision** device has correct correct ownership configured after installing **linux-sgx**, and you can proceed with TDX attestation without rebooting the host.
Jira:RHEL-110112

VM migration no longer fails when using vTPM on shared storage

Before this update, when a virtual Trusted Platform Module (vTPM) data directory was stored on a shared file system, such as NFS, the system failed to create the directory on the destination host during migration, even if it did not exist. This caused virtual machine (VM) migrations to fail. With this update, the system correctly identifies missing vTPM data directories on the destination host and creates them as needed. As a result, virtual machines with a vTPM on shared storage now migrate successfully.

Jira:RHEL-132534[1]

Live VM memory dumps and VM snapshots now work correctly on IBM Z

Previously, attempting to create a memory dump of a running VM by using the **virsh dump --live**

command on an IBM Z host sometimes caused the VM to become unresponsive. In rare cases, creating a snapshot of a running VM can also caused the VM to become unresponsive. With this update, this issue has been fixed, and VMs on IBM Z work as expected in the described scenarios.

Jira:RHELDOCS-21707[1]

12.14. SUPPORTABILITY

Review issues that have been fixed for supportability in Red Hat Enterprise Linux 10.2.

The **rhsm.service** service is running after the**sos** report execution

Before this update, the **sos** report inadvertently started **rhsm.service** service even when it was stopped. This caused the service to run in scenarios where there was no internet connection, generating error messages.
With this fix, the **sos** report no longer starts **rhsm.service** service when it is disabled, improving system stability in offline environments.

Jira:RHEL-112563

Scrub non-alphanumeric passwords are available in the installer logs

Before this update, password detection was strict for obfuscating non-alphanumeric characters. With this release, password scrubbing now accepts non-alphanumeric characters. As a result, password detection no longer rejects non-alphanumeric characters, improving password input flexibility.
Jira:RHEL-121515

Improved IPv6 obfuscation for data privacy

Before this update, the netmask portion of IPv6 addresses remained visible during the data cleaning process. With this release, both the address and the netmask are properly obfuscated, preventing the accidental exposure of network topology.
Jira:RHEL-121517

The **obfuscate_file** function correctly scrubs file content

Before this update, the **obfuscate_file** function overwrote the file content with the filename, causing issues with the main archive population in the cleaner. Consequently, incorrectly overwritten file content in **sos** caused user data corruption. This update introduces the following notable enhancements:

- The **obfuscate_file** function cleans the file content instead of the filename.
- The cleaner's **main_archive** is populated by the parsers first to ensure data integrity.
- The **obfuscate_file** function does not require **short_name**. It uses an implicit value that the cleaner automatically processes.

Jira:RHEL-121531

Enhanced post processing obfuscation in OpenStack Nova

Before this update, the passwords were never scrubbed. With this update, the obfuscation is applied only to the **/var/lib/openstack/config/nova** directory and obfuscating passwords from transport URLs, not the entire URL.
Jira:RHEL-121534

Improper scrubbing fixed in **aap_containerized** to secure passwords

Before this update, the unscrubbed passwords were collected from containerized AAP deployments because of the improper scrubbing in the **aap_containerized** plugin. As a consequence, a password leak occurred in these deployments.
With this release, secret obfuscation has been added to the plugin. As a result, sensitive data is properly obfuscated in the containerized AAP deployments, reducing the risk of password leaks.

Jira:RHEL-142618

12.15. CONTAINERS

Review issues that have been fixed for containers in Red Hat Enterprise Linux 10.2.

Skopeo switches to Sequoia-PGP for OpenPGP signatures in RHEL 10

With this update, Skopeo supports a Sequoia-PGP-based backend for OpenPGP image signatures. Previously, skopeo used **GnuPG** (**gpgme**/**pgpme** bindings) for **OpenPGP** operations. This update includes the following enhancements:

- Verification: the back end is switched from GnuPG to Sequoia-PGP.
- Signing: the current GnuPG workflows continue to exist. New **--sign-by-sq-fingerprint** option allow you to use Sequoia and Sequoia-available keys. Current GnuPG workflows remain supported.
- Algorithm support: Supports modern and post-quantum capable algorithms such as ML-DSA-87+Ed448.
- Improved Skopeo compatibility with FIPS certification.

Jira:RHEL-56364[1]

Buildah and Podman no longer request multiple tokens per operation

Previously, the Buildah and Podman utilities repeatedly requested tokens during each operation. This sometimes caused a race condition in the hosted repository manager.
This update fixes the issue, it prevents multiple token requests which improves the performance and stability of the hosted repository manager.

Jira:RHEL-164030

12.16. RHEL LIGHTSPEED

Review issues that have been fixed for RHEL Lightspeed in Red Hat Enterprise Linux 10.2.

The **lightspeed** keyword is added to **dnf** search metadata for the CLA package

Before this update, the **lightspeed** keyword was missing from the command-line assistant (CLA)

package summary. As a consequence, users could not easily find the package when performing a **dnf** search. With this update, the keyword is added to the package metadata. As a result, users can now find the package by searching for **lightspeed**, which makes the CLA easier to install.
Jira:RHEL-114376

CHAPTER 13. AVAILABLE BPF FEATURES

Review Berkeley Packet Filter (BPF) features available in the Red Hat Enterprise Linux 10.2 kernel to understand your system's capabilities.

Table 13.1. System configuration and other options

Option	Value
unprivileged_bpf_disabled	2 (bpf() syscall restricted to privileged users, admin can change)
JIT enable	1 (enabled)
JIT harden	1 (enabled for unprivileged users)
JIT kallsyms	1 (enabled for root)
Memory limit for JIT for unprivileged users	69267617742848
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
CONFIG_BPF_EVENTS	y
CONFIG_KPROBE_EVENTS	y
CONFIG_UPROBE_EVENTS	y

Option	Value
CONFIG_TRACING	y
CONFIG_FTRACE_SYSCALLS	y
CONFIG_FUNCTION_ERROR_INJECTION	n
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
CONFIG_TEST_BPF	m
CONFIG_HZ	100
bpf() syscall	available

Option	Value
Large insn size limit	available
Bounded loop support	available
ISA extension v2	available
ISA extension v3	available
ISA extension v4	available

Table 13.2. 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	

Program type	Available helpers
lsm	
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 13.3. 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-18041, Jira:RHEL-58682, Jira:RHEL-64019, Jira:RHEL-76835, Jira:RHEL-86312, Jira:RHEL-86320, Jira:RHEL-86534, Jira:RHEL-89601, Jira:RHEL-99331, Jira:RHEL-106502, Jira:RHEL-106849, Jira:RHEL-107003, Jira:RHEL-109113, Jira:RHEL-110192, Jira:RHEL-111220, Jira:RHEL-117520, Jira:RHEL-121208, Jira:RHEL-122674, Jira:RHEL-123220, Jira:RHEL-123663, Jira:RHEL-123664, Jira:RHEL-124694, Jira:RHEL-128906, Jira:RHEL-129675, Jira:RHEL-133085, Jira:RHEL-138729, Jira:RHEL-139826, Jira:RHEL-25071
NetworkManager	Jira:RHEL-95844, Jira:RHEL-100768, Jira:RHEL-108454, Jira:RHEL-109853, Jira:RHEL-5852
NetworkManager-libreswan	Jira:RHEL-33712, Jira:RHEL-67307
Release Notes	Jira:RHELDOCS-21869, Jira:RHELDOCS-19072, Jira:RHELDOCS-19891, Jira:RHELDOCS-19968, Jira:RHELDOCS-20041, Jira:RHELDOCS-20042, Jira:RHELDOCS-20043, Jira:RHELDOCS-19635, Jira:RHELDOCS-21587, Jira:RHELDOCS-20056, Jira:RHELDOCS-20440, Jira:RHELDOCS-20690, Jira:RHELDOCS-18080, Jira:RHELDOCS-19607, Jira:RHELDOCS-18674, Jira:RHELDOCS-18672, Jira:RHELDOCS-18450, Jira:RHELDOCS-20688, Jira:RHELDOCS-20689, Jira:RHELDOCS-20691, Jira:RHELDOCS-20692, Jira:RHELDOCS-20693, Jira:RHELDOCS-20147, Jira:RHELDOCS-20610, Jira:RHELDOCS-21532, Jira:RHELDOCS-16612, Jira:RHELDOCS-21618, Jira:RHELDOCS-21374, Jira:RHELDOCS-19015, Jira:RHELDOCS-19603, Jira:RHELDOCS-19770, Jira:RHELDOCS-19853, Jira:RHELDOCS-20686, Jira:RHELDOCS-20687
adcli	Jira:RHEL-2518, Jira:RHEL-5044, Jira:RHEL-5050, Jira:RHEL-16141
aide	Jira:RHEL-1383
anaconda	Jira:RHEL-16168, Jira:RHEL-58215, Jira:RHEL-58828, Jira:RHEL-66155, Jira:RHEL-95061, Jira:RHEL-95062, Jira:RHEL-96216, Jira:RHEL-135116, Jira:RHEL-151547, Jira:RHEL-147437, Jira:RHEL-80672, Jira:RHEL-67865, Jira:RHEL-74504, Jira:RHEL-69400, Jira:RHEL-83577, Jira:RHEL-58827, Jira:RHEL-58829
ansible-collection-microsoft-sql	Jira:RHEL-69315
ansible-collection-redhat-leapp	Jira:RHEL-141757

Component	Tickets
ansible-core	Jira:RHEL-126844, Jira:RHEL-86829
ansible-freeipa	Jira:RHEL-139147, Jira:RHEL-139258, Jira:RHEL-140606
boost	Jira:RHEL-124169
bootc-image-builder-container	Jira:RHEL-34807
buildah	Jira:RHEL-114411, Jira:RHEL-127903, Jira:RHEL-164030
ca-certificates	Jira:RHEL-120696
capnproto	Jira:RHEL-114452, Jira:RHEL-127899
cepces	Jira:RHEL-121729
chrony	Jira:RHEL-112593
clevis	Jira:RHEL-132188
clevis-pin-tpm2	Jira:RHEL-138591
clevis-pin-trustee	Jira:RHEL-139808
cloud-init	Jira:RHEL-82209, Jira:RHEL-82210
cockpit	Jira:RHEL-112867, Jira:RHEL-4032
cockpit-machines	Jira:RHEL-31993
command-line-assistant	Jira:RHEL-114376
container-selinux	Jira:RHEL-111947
container-tools	Jira:RHEL-127903, Jira:RHEL-67860
coreutils	Jira:RHEL-74146
crash	Jira:RHEL-114659
crun	Jira:RHEL-114419, Jira:RHEL-127903, Jira:RHEL-161090

Component	Tickets
crypto-policies	Jira:RHEL-93769, Jira:RHEL-125889, Jira:RHEL-133522, Jira:RHEL-148560, Jira:RHEL-99813, Jira:RHEL-64746, Jira:RHEL-112392
cups-filters	Jira:RHEL-93944
device-mapper-multipath	Jira:RHEL-118720, Jira:RHEL-133815, Jira:RHEL-141287
distribution	Jira:RHEL-114452, Jira:RHEL-133550, Jira:RHEL-139533, Jira:RHEL-139808, Jira:RHEL-141757, Jira:RHEL-120788, Jira:RHEL-73770
dnf	Jira:RHEL-94331, Jira:RHEL-112730, Jira:RHEL-128445
dnf-plugins-core	Jira:RHEL-94828
dogtag-pki	Jira:RHEL-143038
edk2	Jira:RHEL-66234, Jira:RHEL-68418, Jira:RHEL-82685
elfutils	Jira:RHEL-121665
environment-modules	Jira:RHEL-132336
fapolicyd	Jira:RHEL-1368, Jira:RHEL-94786, Jira:RHEL-118362, Jira:RHEL-131723, Jira:RHEL-114562
firewalld	Jira:RHEL-70357
frr	Jira:RHEL-118620
fuse-overlayfs	Jira:RHEL-128521
fwupd	Jira:RHEL-110760
gcc	Jira:RHEL-105464
gcc-toolset-15	Jira:RHEL-88743
gdm	Jira:RHEL-14524
glibc	Jira:RHEL-65838, Jira:RHEL-102553, Jira:RHEL-111115, Jira:RHEL-111117, Jira:RHEL-114265, Jira:RHEL-115823, Jira:RHEL-118273, Jira:RHEL-119392, Jira:RHEL-119434, Jira:RHEL-137184, Jira:RHEL-140103, Jira:RHEL-139419, Jira:RHEL-146428, Jira:RHEL-150270, Jira:RHEL-142675

Component	Tickets
gnome-control-center / Displays panel	Jira:RHEL-144935
gnome-shell	Jira:RHEL-11918
gnutls	Jira:RHEL-102992
goose	Jira:RSPEED-2846
greenboot-rs	Jira:RHEL-141567
grub2	Jira:RHEL-24510, Jira:RHEL-119685
gvisor-tap-vsock	Jira:RHEL-111948
httpd	Jira:RHEL-145713
ipa	Jira:RHEL-86030, Jira:RHEL-90121, Jira:RHEL-110204, Jira:RHEL-113778, Jira:RHEL-119481, Jira:RHEL-120956, Jira:RHEL-126761, Jira:RHEL-67912, Jira:RHEL-147173, Jira:RHEL-12154
ipmitool	Jira:RHEL-112449
iproute	Jira:RHEL-98263, Jira:RHEL-131660
java-25-openjdk	Jira:RHEL-128409
kdump-utils	Jira:RHEL-29037
kernel	Jira:RHELPLAN-114103
kernel / BPF	Jira:RHEL-78204
kernel / Core	Jira:RHEL-83442, Jira:RHEL-83042
kernel / Crypto	Jira:RHEL-94928, Jira:RHEL-95628, Jira:RHEL-106909
kernel / Debugging-Tracing / EDAC-HERM	Jira:RHEL-45084
kernel / Debugging-Tracing / Ftrace	Jira:RHEL-87219, Jira:RHEL-105766

Component	Tickets
kernel / Debugging-Tracing / Perf	Jira:RHEL-23107, Jira:RHEL-45066, Jira:RHEL-68347, Jira:RHEL-78200, Jira:RHEL-78292
kernel / Debugging-Tracing / rtla	Jira:RHEL-89807
kernel / Desktop / Graphics	Jira:RHEL-88668
kernel / File Systems / VFS	Jira:RHEL-44601
kernel / Networking	Jira:RHEL-77189, Jira:RHEL-100942, Jira:RHEL-130475, Jira:RHEL-130765
kernel / Networking / NIC Drivers	Jira:RHEL-100066, Jira:RHEL-134991, Jira:RHEL-114861, Jira:RHEL-56981
kernel / Networking / Protocol / tcp	Jira:RHEL-115393
kernel / Networking / Protocol / udp	Jira:RHEL-142435
kernel / Networking / Wifi	Jira:RHEL-111098, Jira:RHEL-141347
kernel / Other	Jira:RHEL-2588, Jira:RHEL-65347
kernel / Storage	Jira:RHEL-137767
kernel / Storage / Storage Drivers	Jira:RHEL-115965
kernel / Virtualization	Jira:RHEL-76477
kernel / Virtualization / ESXi	Jira:RHEL-41133
kernel / Virtualization / Hyper-V	Jira:RHEL-75576
kernel / Virtualization / KVM	Jira:RHEL-42486, Jira:RHEL-73000, Jira:RHEL-100313, Jira:RHEL-58218, Jira:RHEL-32892, Jira:RHEL-45585, Jira:RHEL-38957
kernel / io_uring	Jira:RHEL-120700

Component	Tickets
keylime	Jira:RHEL-117442, Jira:RHEL-119028, Jira:RHEL-130158, Jira:RHEL-140896
keylime-agent-rust	Jira:RHEL-117122, Jira:RHEL-117441, Jira:RHEL-140897
kpatch	Jira:RHEL-106283
libdnf	Jira:RHEL-128443
libinput	Jira:RHEL-136390
librepo	Jira:RHEL-126292
libreswan	Jira:RHEL-5299
libslirp	Jira:RHEL-45147
libsolv	Jira:RHEL-86940
libssh	Jira:RHEL-70825, Jira:RHEL-133421
libvirt	Jira:RHEL-82645, Jira:RHEL-122932, Jira:RHEL-132534, Jira:RHEL-7125
libvirt / General	Jira:RHEL-111863, Jira:RHEL-89426
libvirt / Live Migration	Jira:RHEL-104216
libvirt / Networking	Jira:RHEL-79806
libvirt / Storage	Jira:RHEL-80679, Jira:RHEL-118671, Jira:RHEL-135115
linux-sgx	Jira:RHEL-110112, Jira:RHEL-121612
llvm	Jira:RHEL-100887
lvm2	Jira:RHEL-60931
mariadb11.8	Jira:RHEL-115468, Jira:RHEL-171580
mesa	Jira:RHEL-45898
mutter	Jira:RHEL-69291

Component	Tickets
nftables	Jira:RHEL-108861, Jira:RHEL-121194
nginx	Jira:RHEL-33742
nmstate	Jira:RHEL-26350, Jira:RHEL-90096, Jira:RHEL-114959
nodejs	Jira:RHEL-90826
nodejs24	Jira:RHEL-90826
nss	Jira:RHEL-114443
nvme-cli	Jira:RHEL-135994
opencryptoki	Jira:RHEL-100058
openscap	Jira:RHEL-133978
openssh	Jira:RHEL-5281, Jira:RHEL-70824, Jira:RHEL-91181, Jira:RHEL-101440, Jira:RHEL-118406, Jira:RHEL-125929
openssl	Jira:RHEL-45704, Jira:RHEL-72719
openwsman	Jira:RHEL-99191
osbuild	Jira:RHEL-4644
p11-kit	Jira:RHEL-89706, Jira:RHEL-97770, Jira:RHEL-139074
pacemaker	Jira:RHEL-23082, Jira:RHEL-62722
pam	Jira:RHEL-130871
papers	Jira:RHEL-86193
pcp	Jira:RHEL-83866, Jira:RHEL-85456, Jira:RHEL-85457, Jira:RHEL-85725, Jira:RHEL-104669, Jira:RHEL-124897
pcs	Jira:RHEL-7670, Jira:RHEL-7673, Jira:RHEL-76157, Jira:RHEL-76162, Jira:RHEL-84120, Jira:RHEL-111451, Jira:RHEL-126839
php	Jira:RHEL-105827

Component	Tickets
php8.4	Jira:RHEL-105827
pkcs11-provider	Jira:RHEL-68621
plymouth	Jira:RHEL-60198
podman	Jira:RHEL-56365, Jira:RHEL-126644, Jira:RHEL-127903, Jira:RHEL-88122, Jira:RHEL-32266, Jira:RHEL-70218, Jira:RHEL-89373, Jira:RHEL-40641
policycoreutils	Jira:RHEL-94827, Jira:RHEL-111505
postgresql18	Jira:RHEL-116546
pykickstart	Jira:RHEL-34829
python-blivet	Jira:RHEL-122305, Jira:RHEL-53719, Jira:RHEL-158237
python-podman	Jira:RHEL-114423
python3.14	Jira:RHEL-120788
qemu-kvm	Jira:RHEL-71834, Jira:RHEL-81999, Jira:RHEL-58928, Jira:RHEL-66229
qemu-kvm / Devices / CPU Models	Jira:RHEL-28971
qemu-kvm / Devices / Machine Types	Jira:RHEL-104009
qemu-kvm / General	Jira:RHEL-73008
qemu-kvm / Networking	Jira:RHEL-45624
qemu-kvm / Storage	Jira:RHEL-66064
qemu-kvm / Storage / qcow2	Jira:RHEL-151317
rear	Jira:RHEL-84286
resource-agents	Jira:RHEL-115495, Jira:RHEL-116152
rhel-bootc-container	Jira:RHEL-82380, Jira:RHEL-34859

Component	Tickets
rhel-system-roles	Jira:RHEL-46226, Jira:RHEL-46227, Jira:RHEL-66738, Jira:RHEL-104931, Jira:RHEL-110865, Jira:RHEL-114467, Jira:RHEL-115033, Jira:RHEL-123016, Jira:RHEL-123026, Jira:RHEL-123523, Jira:RHEL-127971, Jira:RHEL-128428, Jira:RHEL-129309, Jira:RHEL-136597, Jira:RHEL-136607, Jira:RHEL-137261, Jira:RHEL-138277, Jira:RHEL-144495, Jira:RHEL-144914, Jira:RHEL-145214, Jira:RHEL-145219, Jira:RHEL-145247, Jira:RHEL-150780, Jira:RHEL-150788, Jira:RHEL-151437, Jira:RHEL-73440
rpm	Jira:RHEL-112394, Jira:RHEL-118365, Jira:RHEL-56363
rteval	Jira:RHEL-114927
ruby4.0	Jira:RHEL-133550
rust	Jira:RHEL-111845
rust-podman-sequoia	Jira:RHEL-126677
rust-rpm-sequoia	Jira:RHEL-130960, Jira:RHEL-144414, Jira:RHEL-152461
rust-sequoia-sq	Jira:RHEL-85985
samba	Jira:RHEL-114545
scap-security-guide	Jira:RHEL-152059
selinux-policy	Jira:RHEL-50299, Jira:RHEL-106998, Jira:RHEL-107038, Jira:RHEL-107732, Jira:RHEL-129839, Jira:RHEL-139385, Jira:RHEL-77808
setools	Jira:RHEL-115363
shim	Jira:RHEL-144033
skopeo	Jira:RHEL-127903, Jira:RHEL-56364
snapm	Jira:RHEL-137376
sos	Jira:RHEL-103783, Jira:RHEL-112563, Jira:RHEL-114887, Jira:RHEL-121524, Jira:RHEL-142619, Jira:RHEL-140738, Jira:RHEL-121515, Jira:RHEL-121517, Jira:RHEL-121531, Jira:RHEL-121534, Jira:RHEL-142618
sscg	Jira:RHEL-123675
sssd	Jira:RHEL-4990, Jira:RHEL-11913, Jira:RHEL-94545, Jira:RHEL-127792

Component	Tickets
stratisd	Jira:RHEL-125937
sudo	Jira:RHEL-112100
systemd	Jira:RHEL-109832, Jira:RHEL-92781
systemtap	Jira:RHEL-121663
tbb	Jira:RHEL-33633
tpm2-tools	Jira:RHEL-94930
trustee-guest-components	Jira:RHEL-73770
tuna	Jira:RHEL-116084
tuned	Jira:RHEL-79913
unbound	Jira:RHEL-147790
valgrind	Jira:RHEL-120966
vdo	Jira:RHEL-129906
vim	Jira:RHEL-145868
virt-v2v	Jira:RHEL-13340
virtio-win	Jira:RHEL-91040, Jira:RHEL-1609
virtio-win / user-mode	Jira:RHEL-91041
virtio-win / virtio-win-prewhql	Jira:RHEL-1084, Jira:RHEL-53962, Jira:RHEL-12118
virtiofsd	Jira:RHEL-99895
volume_key	Jira:RHEL-146218
xdp-tools	Jira:RHEL-105793

Component	Tickets
other	Jira:RHELDOCS-21216, Jira:RHELDOCS-19587, Jira:RHELDOCS-22010, Jira:RHELDOCS-18977, Jira:RHELDOCS-21813, Jira:RHELDOCS-21814, Jira:RHELDOCS-21815, Jira:RHELDOCS-21852, Jira:RHELDOCS-21885, Jira:RHELDOCS-21394, Jira:RHELDOCS-20708, Jira:RHELDOCS-21383, Jira:RHELDOCS-21869, Jira:RHELDOCS-21963, Jira:RHELDOCS-20631, Jira:RHELDOCS-21707, Jira:RHELDOCS-19757, Jira:RHELDOCS-20426, Jira:RHELDOCS-20258, Jira:RHELDOCS-20354, Jira:RHELDOCS-16800, Jira:RHELDOCS-19891, Jira:RHELDOCS-19968, Jira:RHELDOCS-20041, Jira:RHELDOCS-20042, Jira:RHELDOCS-20043, Jira:RHELDOCS-20080, Jira:RHELDOCS-19635, Jira:RHELDOCS-21395, Jira:RHELDOCS-22164, Jira:RHELDOCS-21153, Jira:RHELDOCS-22088, Jira:RHELDOCS-22154, Jira:RHEL-97489, Jira:RHELDOCS-20138, Jira:RHELDOCS-18700, Jira:RHELDOCS-18903, Jira:RHELDOCS-18904, Jira:RHELDOCS-18491, Jira:RHELDOCS-18672, Jira:RHELDOCS-18450, Jira:RHELDOCS-20147, Jira:RHELDOCS-20283, Jira:RHELDOCS-20610, Jira:RHELDOCS-16612, Jira:RHELDOCS-22157, Jira:RHELDOCS-21618, Jira:RHELDOCS-21758, Jira:RHELDOCS-21325, Jira:RHELDOCS-21726, Jira:RHELDOCS-19015, Jira:RHELDOCS-19603, Jira:RHELDOCS-18471, Jira:RHELDOCS-19770, Jira:RHELDOCS-19539, Jira:RHELDOCS-19734, Jira:RHELDOCS-19948, Jira:RHELDOCS-19496, Jira:RHELDOCS-19945

APPENDIX B. REVISION HISTORY

Review the revision history to track updates to the Red Hat Enterprise Linux 10.2 Release Notes.

0.0-0

Wed 20 May 2026, Valentina Ashirova (vaashiro@redhat.com)

- Release of the Red Hat Enterprise Linux 10.2 Release Notes.