Red Hat Enterprise Linux 9.4

9.4 Release Notes

Release Notes for Red Hat Enterprise Linux 9.4

Red Hat Enterprise Linux 9.4 9.4 Release Notes

Release Notes for Red Hat Enterprise Linux 9.4

Legal Notice

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 9.4 and document known problems in this release, as well as notable bug fixes, Technology Previews, deprecated functionality, and other details. For information about installing Red Hat Enterprise Linux, see Installation.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE

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

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

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

Submitting feedback through Jira (account required)

1. Log in to the Jira website.

2. Click **Create** in the top navigation bar

3. Enter a descriptive title in the **Summary** field.

4. Enter your suggestion for improvement in the **Description** field. Include links to the relevant parts of the documentation.

5. Click **Create** at the bottom of the dialogue.

CHAPTER 1. OVERVIEW

1.1. MAJOR CHANGES IN RHEL 9.4

Installer and image creation
Key highlights for RHEL image builder:

- From the RHEL 9.4 release distribution and onwards, you can specify arbitrary custom mount points, except for specific paths that are reserved for the operating system.

- You can create different partitioning modes, such as **auto-lvm**, **lvm**, and **raw**.

- You can customize tailoring options for a profile and add it to your blueprint customizations by using selected and unselected options, to add and remove rules.

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

RHEL for Edge
Key highlights for RHEL for Edge:

- You can now create FIPS compliant RHEL for Edge images.

- With this Technology Preview, you can now use the FDO onboarding process by storing and querying Owner Vouchers from the Sqlite or Postgresql databases.

For more information, see New features - RHEL for Edge .

Security
The **SELinux** userspace release 3.6 introduces deny rules for further customizing SELinux policies.

The **Keylime** server components, the verifier and registrar, are available as containers.

The **Rsyslog** log processing system introduces customizable TLS/SSL encryption settings and additional options that relate to capability dropping.

The **OpenSSL** TLS toolkit adds a drop-in directory for provider-specific configuration files.

The Linux kernel cryptographic API (**libkcapi**) 1.4.0 introduces new tools and options. Notably, with the new **-T** option, you can specify target file names in hash-sum calculations.

The **stunnel** TLS/SSL tunneling service 5.71 changes the behavior of OpenSSL 1.1 and later versions in FIPS mode. Besides this change, version 5.71 provides many new features such as support for modern PostgreSQL clients.

See New features - Security for more information.

Dynamic programming languages, web and database servers
Later versions of the following Application Streams are now available:

- **Python 3.12**

- **Ruby 3.3**

- **PHP 8.2**

- **nginx 1.24**

- MariaDB 10.11

- PostgreSQL 16

The following components have been upgraded:

- **Git** to version 2.43.0

- **Git LFS** to version 3.4.1

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

Compilers and development tools

Updated performance tools and debuggers

The following performance tools and debuggers have been updated in RHEL 9.4:

- Valgrind 3.22

- SystemTap 5.0

- elfutils 0.190

Updated performance monitoring tools

The following performance monitoring tools have been updated in RHEL 9.4:

- PCP 6.2.0

Updated compiler toolsets

The following compiler toolsets have been updated in RHEL 9.4:

- GCC Toolset 13

- LLVM Toolset 17.0.6

- Rust Toolset 1.75.1

- Go Toolset 1.21.7

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

Identity Management

Key highlights for Identity Management:

- You can enable and configure passwordless authentication in SSSD to use a biometric device that is compatible with the FIDO2 specification, for example a YubiKey.

See New Features – Identity Management for more information.

Virtualization

RHEL 9.4 introduces full support for KVM virtual machines on the 64-bit ARM architecture.

In addition, external snapshot for virtual machines are now fully supported and have become the default mechanism for a number of snapshot operations.

For more information about virtualization features introduced in this release, see New features – Virtualization.

Containers

Notable changes include:

- The **podman farm build** command for creating multi-architecture container images is available as a Technology Preview.

- Podman now supports **containers.conf** modules to load a predetermined set of configurations.

- The Container Tools packages have been updated.

- Podman v4.9 RESTful API now displays data of progress when you pull or push an image to the registry.

- SQLite is now fully supported as a default database backend for Podman.

- **Containerfile** now supports multi-line HereDoc instructions.

- **pasta** as a network name has been deprecated.

- The BoltDB database backend has been deprecated.

- The **container-tools:4.0** module has been deprecated.

- The Container Network Interface (CNI) network stack is deprecated and will be removed in a future release.

See New features - Containers for more information.

1.2. IN-PLACE UPGRADE

In-place upgrade from RHEL 8 to RHEL 9

The supported in-place upgrade paths currently are:

- From RHEL 8.8 to RHEL 9.2, and RHEL 8.10 to RHEL 9.4 on the following architectures:

  - 64-bit Intel

  - 64-bit AMD

  - 64-bit ARM

  - IBM POWER 9 (little endian) and later

  - IBM Z architectures, excluding z13

- From RHEL 8.6 to RHEL 9.0 and RHEL 8.8 to RHEL 9.2 on systems with SAP HANA

For more information, see Supported in-place upgrade paths for Red Hat Enterprise Linux .

For instructions on performing an in-place upgrade, see Upgrading from RHEL 8 to RHEL 9 .

If you are upgrading to RHEL 9.2 with SAP HANA, ensure that the system is certified for SAP before the upgrade. For instructions on performing an in-place upgrade on systems with SAP environments, see How to in-place upgrade SAP environments from RHEL 8 to RHEL 9 .

Notable enhancements include:

- New logic has been implemented to determine the expected states of the **systemd** services after the upgrade.

- Locally stored DNF repositories can now be used for the in-place upgrade.

- You can now configure DNF to be able to upgrade by using proxy.

- Issues with performing the in-place upgrade with custom DNF repositories accessed by using HTTPS have been fixed.

- If the **/etc/pki/tls/openssl.cnf** configuration file has been modified, the file is now replaced with the target default OpenSSL configuration file during the upgrade to prevent issues after the upgrade. See the pre-upgrade report for more information.

In-place upgrade from RHEL 7 to RHEL 9

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

1.3. RED HAT CUSTOMER PORTAL LABS

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

- Registration Assistant

- Kickstart Generator

- Red Hat Product Certificates

- Red Hat CVE Checker

- Kernel Oops Analyzer

- Red Hat Code Browser

- VNC Configurator

- Red Hat OpenShift Container Platform Update Graph

- Red Hat Satellite Upgrade Helper

- JVM Options Configuration Tool

- Load Balancer Configuration Tool

- Red Hat OpenShift Data Foundation Supportability and Interoperability Checker

- Ansible Automation Platform Upgrade Assistant

- Ceph Placement Groups (PGs) per Pool Calculator

- Yum Repository Configuration Helper

1.4. ADDITIONAL RESOURCES

Capabilities and limits of Red Hat Enterprise Linux 9 as compared to other versions of the system are available in the Knowledgebase article Red Hat Enterprise Linux technology capabilities and limits .

Information regarding the Red Hat Enterprise Linux **life cycle** is provided in the Red Hat Enterprise Linux Life Cycle document.

The Package manifest document provides a **package listing** for RHEL 9, including licenses and application compatibility levels.

Application compatibility levels are explained in the Red Hat Enterprise Linux 9: Application Compatibility Guide document.

Major **differences between RHEL 8 and RHEL 9**, including removed functionality, are documented in Considerations in adopting RHEL 9 .

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

The **Red Hat Insights** service, which enables you to proactively identify, examine, and resolve known technical issues, is available with all RHEL subscriptions. For instructions on how to install the Red Hat Insights client and register your system to the service, see the Red Hat Insights Get Started page.

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

---

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

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 9.4 is distributed with the kernel version 5.14.0-427.13.1, which provides support for the following architectures at the minimum required version (stated in parentheses):

- AMD and Intel 64-bit architectures (x86-64-v2)

- The 64-bit ARM architecture (ARMv8.0-A)

- IBM Power Systems, Little Endian (POWER9)

- 64-bit IBM Z (z14)

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

CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 9

3.1. INSTALLATION

Red Hat Enterprise Linux 9 is installed using ISO images. Two types of ISO image are available for the AMD64, Intel 64-bit, 64-bit ARM, IBM Power Systems, and IBM Z architectures:

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

  > **NOTE**
  >
  > The Installation ISO image is in multiple GB size, and as a result, it might not fit on optical media formats. A USB key or USB hard drive is recommended when using the Installation ISO image to create bootable installation media. You can also use the Image Builder tool to create customized RHEL images. For more information about Image Builder, see the *Composing a customized RHEL system image* document.

- Boot ISO: A minimal boot ISO image that is used to boot into the installation program. This option requires access to the BaseOS and AppStream repositories to install software packages. The repositories are part of the Installation ISO image. You can also register to Red Hat CDN or Satellite during the installation to use the latest BaseOS and AppStream content from Red Hat CDN or Satellite.

See the Performing a standard RHEL 9 installation document for instructions on downloading ISO images, creating installation media, and completing a RHEL installation. For automated Kickstart installations and other advanced topics, see the Performing an advanced RHEL 9 installation document.

3.2. REPOSITORIES

Red Hat Enterprise Linux 9 is distributed through two main repositories:

- BaseOS

- AppStream

Both repositories are required for a basic RHEL installation, and are available with all RHEL subscriptions.

Content in the BaseOS repository is intended to provide the core set of the underlying operating system functionality that provides the foundation for all installations. This content is available in the RPM format and is subject to support terms similar to those in previous releases of RHEL. For more information, see the Scope of Coverage Details document.

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

In addition, the CodeReady Linux Builder repository is available with all RHEL subscriptions. It provides additional packages for use by developers. Packages included in the CodeReady Linux Builder repository are unsupported.

For more information about RHEL 9 repositories and the packages they provide, see the Package manifest.

3.3. APPLICATION STREAMS

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

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

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

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

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

Some additional Application Stream versions will be distributed as modules with a shorter life cycle in future minor RHEL 9 releases. Modules are collections of packages representing a logical unit: an application, a language stack, a database, or a set of tools. These packages are built, tested, and released together.

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

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

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

3.4. PACKAGE MANAGEMENT WITH YUM/DNF

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

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

For more information, see Managing software with the DNF tool.

CHAPTER 4. NEW FEATURES

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 9.4.

4.1. INSTALLER AND IMAGE CREATION

Support to add customized files for SCAP security profile to a blueprint

With this enhancement, you can now add customized tailoring options for a profile to the **osbuild-composer** blueprint customizations by using the following options:

- **selected** for the list of rules that you want to add

- **unselected** for the list of rules that you want to remove

With the default **org.ssgproject.content** rule namespace, you can omit the prefix for rules under this namespace. For example: the **org.ssgproject.content_grub2_password** and **grub2_password** are functionally equivalent.

When you build an image from that blueprint, it creates a tailoring file with a new tailoring profile ID and saves it to the image as **/usr/share/xml/osbuild-oscap-tailoring/tailoring.xml**. The new profile ID will have **_osbuild_tailoring** appended as a suffix to the base profile. For example, if you use the **cis** base profile, **xccdf_org.ssgproject.content_profile_cis_osbuild_tailoring**.

Jira:RHELDOCS-17792[1]

Minimal RHEL installation now installs only the s390utils-core package

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

Bugzilla:1932480[1]

4.2. SECURITY

Keylime verifier and registrar containers available

You can now configure Keylime server components, the verifier and registrar, as containers. When configured to run inside a container, the Keylime registrar monitors the tenant systems from the container without any binaries on the host. The container deployment provides better isolation, modularity, and reproducibility of Keylime components.

Jira:RHELDOCS-16721[1]

libkcapi now provides an option for specifying target file names in hash-sum calculations

This update of the **libkcapi** (Linux kernel cryptographic API) packages introduces the new option **-T** for specifying target file names in hash-sum calculations. The value of this option overrides file names specified in processed HMAC files. You can use this option only with the **-c** option, for example:

```
$ sha256hmac -c <hmac_file> -T <target_file>
```

Jira:RHEL-15298[1]

Finer control over MACs in SSH with **crypto-policies**

You can now set additional options for message authentication codes (MACs) for the SSH protocol in the system-wide cryptographic policies (**crypto-policies**). With this update, the **crypto-policies** option **ssh_etm** has been converted into a tri-state **etm@SSH** option. The previous **ssh_etm** option has been deprecated.

You can now set **ssh_etm** to one of the following values:

ANY

Allows both **encrypt-then-mac** and **encrypt-and-mac** MACs.

DISABLE_ETM

Disallows **encrypt-then-mac** MACs.

DISABLE_NON_ETM

Disallows MACs that do not use **encrypt-then-mac**.

Note that ciphers that use implicit MACs are always allowed because they use authenticated encryption.

Jira:RHEL-15925

The **semanage fcontext** command no longer reorders local modifications

The **semanage fcontext -l -C** command lists local file context modifications stored in the **file_contexts.local** file. The **restorecon** utility processes the entries in the **file_contexts.local** from the most recent entry to the oldest. Previously, **semanage fcontext -l -C** listed the entries in an incorrect order. This mismatch between processing order and listing order caused problems when managing SELinux rules. With this update, **semanage fcontext -l -C** displays the rules in the correct and expected order, from the oldest to the newest.

Jira:RHEL-24462[1]

Additional services confined in the SELinux policy

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

- **nvme-stas**

- **rust-afterburn**

- **rust-coreos-installer**

- **bootc**

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

Jira:RHEL-12591[1]

New SELinux policy module for the SAP HANA service

This update adds additional rules to the SELinux policy for the SAP HANA service. As a result, the service now runs successfully in SELinux enforcing mode in the **sap_unconfined_t** domain.

Jira:RHEL-21452

The **glusterd** SELinux module moved to a separate **glusterfs-selinux** package

With this update, the **glusterd** SELinux module is maintained in the separate **glusterfs-selinux** package. The module is therefore no longer part of the **selinux-policy** package. For any actions that concern the **glusterd** module, install and use the **glusterfs-selinux** package.

Jira:RHEL-1548

The **fips.so** library for OpenSSL provided as a separate package

OpenSSL uses the **fips.so** shared library as a FIPS provider. With this update, the latest version of **fips.so** submitted to the National Institute of Standards and Technology (NIST) for certification is in a separate package to ensure that future versions of OpenSSL use certified code or code undergoing certification.

Jira:RHEL-23474[1]

The **chronyd-restricted** service is confined by the SELinux policy

This update adds additional rules to the SELinux policy that confine the new **chronyd-restricted** service. As a result, the service now runs successfully in SELinux.

Jira:RHEL-18219

OpenSSL adds a drop-in directory for provider configuration

The OpenSSL TLS toolkit supports provider APIs for installation and configuration of modules that provide cryptographic algorithms. With this update, you can place provider-specific configuration in separate **.conf** files in the **/etc/pki/tls/openssl.d** directory without modifying the main OpenSSL configuration file.

Jira:RHEL-17193

SELinux user-space components rebased to 3.6

The SELinux user-space components **libsepol**, **libselinux**, **libsemanage**, **policycoreutils**, **checkpolicy**, and **mcstrans** library package have been rebased to 3.6. This version provides various bug fixes, optimizations and enhancements, most notably:

- Added support for **deny** rules in CIL.

- Added support for **notself** and **other** keywords in CIL.

- Added the **getpolicyload** binary that prints the number of policy reloads performed on the current system.

Jira:RHEL-16233

GnuTLS rebased to 3.8.3

The GnuTLS package has been rebased to upstream version 3.8.3 This version provides various bug fixes and enhancements, most notably:

- The **gnutls_hkdf_expand** function now accepts only arguments with lengths less than or equal to 255 times hash digest size, to comply with RFC 5869 2.3.

- Length limit for **TLS PSK** usernames has been increased to 65535 characters.

- The **gnutls_session_channel_binding** API function performs additional checks when **GNUTLS_CB_TLS_EXPORTER** is requested accordingly to RFC9622 4.2.

- The **GNUTLS_NO_STATUS_REQUEST** flag and the **%NO_STATUS_REQUEST** priority modifier have been added to allow disabling of the **status_request** TLS extension on the client side.

- GnuTLS now checks the contents of the Change Cipher Spec message to be equal to 1 when the TLS version is older than 1.3.

- ClientHello extensions order is randomized by default.

- GnuTLS now supports EdDSA key generation on PKCS #11 tokens, which previously did not work.

Jira:RHEL-14891[1]

nettle rebased to 3.9.1

The **nettle** library package has been rebased to 3.9.1. This version provides various bug fixes, optimizations and enhancements, most notably:

- Added balloon password hashing

- Added SIV-GCM authenticated encryption mode

- Added Offset Codebook Mode authenticated encryption mode

- Improved performance of the SHA-256 hash function on 64-bit IBM Z, AMD and Intel 64-bit architectures

- Improved performance of the Poly1305 hash function on IBM Power Systems, Little Endian, AMD and Intel 64-bit architectures

Jira:RHEL-14890[1]

p11-kit rebased to 0.25.3

The **p11-kit** packages have been updated to upstream version 0.25.3. The packages contain the **p11-kit** tool for managing PKCS #11 modules, the **trust** tool for operating on the trust policy store, and the **p11-kit** library. Notable enhancements include the following:

- Added support for PKCS #11 version 3.0

- The **pkcs11.h** header file:

  - Added ChaCha20/Salsa20, Poly1305 and IBM-specific mechanisms and attributes

  - Added AES-GCM mechanism parameters for message-based encryption

- The **p11-kit** tool:

  - Added utility commands to list and manage objects of a token (**list-tokens**, **list-mechanisms**, **list-objects**, **import-object**, **export-object**, **delete-object**, and **generate-keypair**)

  - Added utility commands to manage PKCS#11 profiles of a token (**list-profiles**, **add-profile**, and **delete-profile**)

- Added the **print-config** command for printing merged configuration

- The **trust** tool:

  - Added the **check-format** command to validate the format of **.p11-kit** files

Jira:RHEL-14834[1]

libkcapi rebased to 1.4.0

The **libkcapi** library, which provides access to the Linux kernel crypto API, has been rebased to upstream version 1.4.0. The update includes various enhancements and bug fixes, most notably:

- Added the **sm3sum** and **sm3hmac** tools.

- Added the **kcapi_md_sm3** and **kcapi_md_hmac_sm3** APIs.

- Added SM4 convenience functions.

- Fixed support for link-time optimization (LTO).

- Fixed LTO regression testing.

- Fixed support for AEAD encryption of an arbitrary size with **kcapi-enc**.

Jira:RHEL-5367[1]

User and group creation in OpenSSH uses the **sysusers.d** format

Previously, OpenSSH used static **useradd** scripts. With this update, OpenSSH uses the **sysusers.d** format to declare system users, which makes it possible to introspect system users.

Jira:RHEL-5222

OpenSSH limits artificial delays in authentication

OpenSSH's response after login failure is artificially delayed to prevent user enumeration attacks. This update introduces an upper limit on such delays when remote authentication takes too long, for example in privilege access management (PAM) processing.

Jira:RHEL-2469[1]

stunnel rebased to 5.71

The **stunnel** TLS/SSL tunneling service has been rebased to upstream version 5.71.

Notable new features include:

- Added support for modern PostgreSQL clients.

- You can use the **protocolHeader** service-level option to insert custom **connect** protocol negotiation headers.

- You can use the **protocolHost** option to control the client SMTP protocol negotiation HELO/EHLO value.

- Added client-side support for Client-side **protocol = ldap**.

- You can now configure session resumption by using the service-level **sessionResume** option.

- Added support to request client certificates in server mode with **CApath** (previously, only **CAfile** was supported).

- Improved file reading and logging performance.

- Added support for configurable delay for the **retry** option.

- In client mode, OCSP stapling is requested and verified when **verifyChain** is set.

- In server mode, OCSP stapling is always available.

- Inconclusive OCSP verification breaks TLS negotiation. You can disable this by setting **OCSPrequire = no**.

Jira:RHEL-2468[1]

New options for dropping capabilities in Rsyslog

You can now configure Rsyslog's behavior when dropping capabilities by using the following global options:

libcapng.default

　　Determines Rsyslog's actions when it encounters errors while dropping capabilities. The default value is **on**, which caused Rsyslog to exit if an error related to **libcapng-related** occurs.

libcapng.enable

　　Determines whether Rsyslog drops capabilities during startup. If this option is disabled, **libcapng.default** has no impact.

Jira:RHEL-943[1]

audit rebased to 3.1.2

The Linux Audit system has been updated to version 3.1.2, which provides bug fixes, enhancements, and performance improvements over the previously released version 3.0.7. Notable enhancements include:

- The **auparse** library now interprets unnamed and anonymous sockets.

- You can use the new keyword **this-hour** in the **start** and **end** options of the **ausearch** and **aureport** tools.

- Support for the **io_uring** asynchronous I/O API has been added.

- User-friendly keywords for signals have been added to the **auditctl** program.

- Handling of corrupt logs in **auparse** has been improved.

- The **ProtectControlGroups** option is now disabled by default in the **auditd** service.

- Rule checking for the exclude filter has been fixed.

- The interpretation of **OPENAT2** fields has been enhanced.

- The **audispd af_unix** plugin has been moved to a standalone program.

- The Python binding has been changed to prevent setting Audit rules from the Python API. This change was made due to a bug in the Simplified Wrapper and Interface Generator (SWIG).

Jira:RHEL-14896[1]

Rsyslog rebased to 8.2310

The Rsyslog log processing system has been rebased to upstream version 8.2310. This update introduces significant enhancements and bug fixes. Most notable enhancements include:

Customizable TLS/SSL encryption settings

In previous versions, configuring TLS/SSL encryption settings for separate connections was limited to global settings. With the latest version, you can now define unique TLS/SSL settings for each individual connection in Rsyslog. This includes specifying different CA certificates, private keys, public keys, and CRL files for enhanced security and flexibility. For detailed information and usage, see documentation provided in the **rsyslog-doc** package.

Refined capability dropping feature

You can now set additional options that relate to capability dropping. You can disable capability dropping by setting the **libcapng.enable** global option to **off**. For more information, see RHEL-943.

Jira:RHEL-937, Jira:RHEL-943

SCAP Security Guide rebased to 0.1.72

The SCAP Security Guide (SSG) packages have been rebased to upstream version 0.1.72. This version provides bug fixes and various enhancements, most notably:

- CIS profiles are updated to align with the latest benchmarks.

- The PCI DSS profile is aligned with the PCI DSS policy version 4.0.

- STIG profiles are aligned with the latest DISA STIG policies.

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

Jira:RHEL-21425

4.3. RHEL FOR EDGE

Support for building FIPS enabled RHEL for Edge images

This enhancement adds support for building FIPS enabled RHEL for Edge images for the following images types:

- **edge-installer**

- **edge-simplified-installer**

- **edge-raw-image**

- **edge-ami**

- **edge-vsphere**

> **IMPORTANT**
>
> You can enable FIPS mode only during the image provisioning process. You cannot change to FIPS mode after the non-FIPS image build starts.

Jira:RHELDOCS-17263[1]

4.4. SHELLS AND COMMAND-LINE TOOLS

openCryptoki rebased to version 3.22.0

The **opencryptoki** package has been updated to version 3.22.0. Notable changes include:

- Added support for the **AES-XTS** key type by using the **CPACF** protected keys.

- Added support for managing certificate objects.

- Added support for public sessions with the **no-login** option.

- Added support for logging in as the Security Officer (SO).

- Added support for importing and exporting the **Edwards** and **Montgomery** keys.

- Added support for importing the **RSA-PSS** keys and certificates.

- For security reasons, the 2 key parts of an AES-XTS key should not be the same. This update adds checks to the key generation and import process to ensure this.

- Various bug fixes have been implemented.

Jira:RHEL-11412[1]

4.5. INFRASTRUCTURE SERVICES

synce4l rebased to version 1.0.0

The **synce4l** protocol has been updated to version 1.0.0. This update adds support for kernel Digital Phase Locked Loop (DPLL) interface.

Jira:RHEL-10089[1]

chrony rebased to version 4.5

The **chrony** suite has been updated to version 4.5. Notable changes include:

- Added support for the AES-GCM-SIV cipher to shorten Network Time Security (NTS) cookies to improve reliability of NTS over the internet, where some providers block or limit the rate of longer Network Time Protocol (NTP) messages.

- Added periodic refresh of IP addresses of NTP sources specified by hostname. The default interval is two weeks and it can be disabled by adding **refresh 0** parameter to the **chrony.conf** file.

- Improved automatic replacement of unreachable NTP sources.

- Improved logging of important changes made by the **chronyc** utility.

- Improved logging of source selection failures and falsetickers.

- Added the **hwtstimeout** directive to configure timeout for late hardware transmit timestamps.

- Added experimental support for corrections provided by Precision Time Protocol (PTP) transparent clocks to reach accuracy of PTP with hardware timestamping.

- Added the **chronyd-restricted** service as an alternative service for minimal client-only configurations where the **chronyd** service can be started without **root** privileges.

- Fixed the **presend** option in **interleaved** mode.

- Fixed reloading of modified sources specified by IP address from the **sourcedir** directories.

Jira:RHEL-6522

linuxptp rebased to version 4.2

The **linuxptp** protocol has been updated to version 4.2. Notable changes include:

- Added support for multiple domains in the **phc2sys** utility.

- Added support for notifications on clock updates and changes in the Precision Time Protocol (PTP) parent dataset, for example, clock class.

- Added support for PTP Power Profile, namely IEEE C37.238-2011 and IEEE C37.238-2017.

Jira:RHEL-2026

4.6. NETWORKING

The nft utility can now reset nftables rule-contained states

With this enhancement, you can use the **nft reset** command to reset **nftables** rule-contained states. For example, use this feature to reset counter and quota statement values.

Jira:RHEL-5980[1]

Marvell Octeon PCIe Endpoint Network Interface Controller driver is available

This enhancement has added the **octeon_ep** driver. You can use it for networking of Marvell's Octeon PCIe Endpoint network interface cards. The host drivers act as PCI Express (PCIe) endpoint network interface (NIC) to support Marvell OCTEON TX2 CN106XX, a 24 N2 cores Infrastructure Processor Family. By using OCTEON TX2 driver as a PCIe NIC, you can use OCTEON TX2 as a PCIe endpoint in various products: security firewalls, 5G Open Radio Access Network (ORAN) and Virtual RAN (VRAN) applications and data processing offloading applications.

Currently, you can use it with the following devices:

- Network controller: Cavium, Inc. Device b100

- Network controller: Cavium, Inc. Device b200

- Network controller: Cavium, Inc. Device b400

- Network controller: Cavium, Inc. Device b900

- Network controller: Cavium, Inc. Device ba00

- Network controller: Cavium, Inc. Device bc00

- Network controller: Cavium, Inc. Device bd00

Jira:RHEL-9308[1]

NetworkManager now supports configuring the **switchdev** mode for advanced hardware offload

With this enhancement, you can configure the following new properties in NetworkManager connection profiles:

- **sriov.eswitch-mode**

- **sriov.eswitch-inline-mode**

- **sriov.eswitch-encap-mode**

With these properties, you can configure the eSwitch of smart network interface controllers (Smart NICs). For example, use the **sriov.eswitch-mode** setting to change the mode from **legacy SR-IOV** to **switchdev** to use advanced hardware offload features.

Jira:RHEL-1441

NetworkManager supports changing **ethtool** channel settings

A network interface can have multiple interrupt request (IRQs) and associated packet queues called **channels**. With this enhancement, NetworkManager connection profiles can specify the number of channels to assign to an interface through connection properties **ethtool.channels-rx**,**ethtool.channels-tx**,**ethtool.channels-other**, and **ethtool.channels-combined**.

Jira:RHEL-1471[1]

Nmstate can now create a YAML file to revert settings

With this enhancement, Nmstate can create a "revert configuration file" that contains the differences between the current network settings and a YAML file with the new configuration that you want to apply. If the settings do not work as expected after you applied the YAML file, you can use the revert configuration file to restore the previous settings:

1. Create a YAML file, for example, **new.yml** with the configuration that you want to apply.

2. Create a revert configuration file that contains the differences between intended settings in **new.yml** and the current state:

   ```
   # nmstatectl gr new.yml > revert.yml
   ```

3. Apply the configuration from **new.yml**.

4. If you want now to switch back to the previous state, apply **revert.yml**.

Alternatively, you can use the **NetworkState::generate_revert(current)** call if you use the Nmstate API to create a revert configuration.

Jira:RHEL-1434

Nmstate API configures VPN connection based on IPsec configuration

The Libreswan utility is an implementation of IPsec for configuring VPNs. With this update, by using **nmstatectl**, you can configure IPsec-based authentication types along with configuration modes (tunnel and transport) and network layouts (**host-to-subnet**, **host-to-host**, **subnet-to-subnet**).

Jira:RHEL-1605

nmstate now supports the priority bond property

With this update, you can set the priority of bond ports in the **nmstate** framework by using the **priority** property in the **ports-config** section of the configuration file. An example YAML file can look as follows:

```
---
interfaces:
- name: bond99
  type: bond
  state: up
  link-aggregation:
    mode: active-backup
    ports-config:
    - name: eth2
      priority: 15
```

When an active port within the bonded interface is down, the RHEL kernel elects the next active port that has the highest numerical value in the **priority** property from the pool of all backup ports.

The **priority** property is relevant for the following modes of the bond interface:

- **active-backup**

- **balance-tlb**

- **balance-alb**

Jira:RHEL-1438[1]

NetworkManager wifi connections support a new MAC address-based privacy option

With this enhancement, you can configure NetworkManager to associate a random-generated MAC address with the Service Set Identifier (SSID) of a wifi network. This enables you to permanently use a random but consistent MAC address for a wifi network even if you delete a connection profile and recreate it. To use this new feature, set the **802-11-wireless.cloned-mac-address** property of a wifi connection profile to **stable-ssid**.

Jira:RHEL-16470

Introduction of new nmstate attributes for the VLAN interface

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

- **registration-protocol**: VLAN Registration Protocol. The valid values are **gvrp** (GARP VLAN Registration Protocol), **mvrp** (Multiple VLAN Registration Protocol), and **none**.

- **reorder-headers**: reordering of output packet headers. The valid values are **true** and **false**.

- **loose-binding**: loose binding of the interface to the operating state of its primary device. The valid values are **true** and **false**.

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

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

Jira:RHEL-19142

ipv4.dhcp-client-id set to none prevents sending a client-identifier

If the **client-identifier** option is not set in NetworkManager, then the actual value depends on the type of DHCP clients in use, such as NetworkManager **internal** DHCP client or **dhclient**. Generally, DHCP clients send a **client-identifier**. Therefore, in almost all cases, you do not need to set the **none** option. As a result, this option is only useful in case of some unusual DHCP server configurations that require clients to not send a **client-identifier**.

Jira:RHEL-1469

nmstate now supports creating MACsec interfaces

With this update, the users of the **nmstate** framework can configure MACsec interfaces to protect their communication on Layer 2 of the Open Systems Interconnection (OSI) model. As a result, there is no need to encrypt individual services later on Layer 7. Also, the feature eliminates associated challenges such as managing large amounts of certificates for each endpoint.

For more information, see Configuring a MACsec connection using nmstatectl .

Jira:RHEL-1420

netfilter update

The **kernel** package has been upgraded to version 5.14.0-405 in RHEL 9. As a result, the rebase also provided multiple enhancements and bug fixes in the **netfilter** component of the RHEL kernel. The most notable change includes:

- The **nftables** subsystem is able to match various inner header fields of the tunnel packets. This enables more granular and effective control over network traffic, especially in environments where tunneling protocols are used.

Jira:RHEL-16630[1]

firewalld now avoids unnecessary firewall rule flushes

The **firewalld** service does not remove all existing rules from the **iptables** configuration if both following conditions are met:

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

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

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

Jira:RHEL-427[1]

The **ss** utility adds visibility improvement to TCP bound-inactive sockets

The **iproute2** suite provides a collection of utilities to control TCP/IP networking traffic. TCP bound-inactive sockets are attached to an IP address and a port number but neither connected nor listening on TCP ports. The socket services (**ss**) utility adds support for the kernel to dump TCP bound-inactive sockets. You can view those sockets with the following command options:

- **ss --all**: to dump all sockets including TCP bound-inactive ones

- **ss --bound-inactive**: to dump only bound-inactive sockets

Jira:RHEL-21223[1]

The Nmstate API now supports SR-IOV VLAN 802.1ad tagging

With this enhancement, you can now use the Nmstate API to enable hardware-accelerated Single-Root I/O Virtualization (SR-IOV) Virtual Local Area Network (VLAN) 802.1ad tagging on cards whose firmware supports this feature.

Jira:RHEL-1425

The TCP Illinois congestion algorithm kernel module is re-enabled

TCP Illinois is a variant of the TCP protocol. Customers like Internet Service Providers (ISP) experience sub-optimal performance without TCP Illinois algorithm and network traffic does not scale well even when using Bandwidth and Round-trip propagation time (BBR) algorithm that results into high latency. As a result, TCP Illinois algorithm can produce slightly higher average throughput, fairer network resources allocation, and compatibility.

Jira:RHEL-5736[1]

The **iptables** utility rebased to version 1.8.10

The **iptables** utility defines rules for packet filtering to manage firewall. This utility has been rebased. Notable changes include:

Notable features:

- Add support for newer chunk types in **sctp** match

- Align ip6tables opt-in column if empty helps when piping output to **jc --iptables**

- Print numeric protocol numbers with **--numeric** for a more stable output

- More translations for *tables-translate utilities with improved output formatting

- Several manual page improvements

Notable fixes:

- **iptables-restore** error messages incorrectly pointing at the COMMIT line

- Broken **-p Length** match in ebtables

- Broken ebtables among match when used in multiple rules restored through **ebtables-restore**

- Program could crash when renaming a chain depending on the number of chains already present

- Non-critical memory leaks

- Missing broute table support in ebtables after the switch to nft-variants

- Broken ip6tables rule counter setting with '-c' option

- Unexpected error message when listing a non-existent chain

- Potential false-positive ebtables rule comparison if among match is used

- Prohibit renaming a chain to an invalid name

- Stricter checking of "chain lines" in iptables-restore input to detect invalid chain names

- Non-functional built-in chain policy counters

Jira:RHEL-14147

nftables rebased to version 1.0.9

The **nftables** utility has been upgraded to version 1.0.9, which provides multiple bug fixes and enhancements. Notable changes include:

- Improvements to the **--optimize** command option

- Extended the Python nftables class

- Improved behavior when dealing with rules created by **iptables-nft**

- Support accessing fields of vxlan-encapsulated headers

- Initial support for GRE, Geneve, and GRETAP protocols

- New **reset rule(s)** commands to reset rule counters, quotas

- New **destroy** command deletes things only if they exist

- New **last** statement recording when it has seen a packet for the last time

- Add and remove devices from netdev-family chains

- New **meta broute** expression to emulate ebtables' broute functionality

- Fixed miscellaneous memory leaks

- Fixed wrong location in error messages in corner-cases

- Set and map statements missing in JSON output

Jira:RHEL-14191

firewalld rebased to version 1.3

The **firewalld** package has been upgraded to version 1.3, which provides multiple bug fixes and enhancements. Notable changes include:

- New **--reset-defaults** CLI option: This option resets configuration of the **firewalld** service to defaults. This allows users to erase **firewalld** configuration and start over with the default settings.

- Enable the **--add-masquerade** CLI option for policies with **ingress-zone=ZONE**, where **ZONE** has interfaces assigned with the **--add-interface** CLI option. This removes a restriction and enables usage of interfaces (instead of sources) in common scenarios.

The reasons to introduce these features:

- **--reset-defaults** was implemented to reset the firewall to the default configuration.

- Using interfaces allows change of IP address without impacting firewall configuration.

As a result, users can perform the following actions:

- Reset the configuration

- Combine **--add-maquerade** with **--add-interface** while using policies

Jira:RHEL-14485

4.7. KERNEL

Kernel version in RHEL 9.4

Red Hat Enterprise Linux 9.4 is distributed with the kernel version 5.14.0-427.13.1.

rteval now supports adding and removing arbitrary CPUs from the default measurement CPU list

With the **rteval** utility, you can add (using the + sign) or subtract (using the - sign) CPUs to the default measurement CPU list when using the **--measurement-cpulist** parameter, instead of having to specify an entire new list. Additionally, **--measurement-run-on-isolcpus** is introduced for adding the set of all isolated CPUs to the default measurement CPU list. This option covers the most common use case of a real-time application running on isolated CPUs. Other use cases require a more generic feature. For example, some real-time applications used one isolated CPU for housekeeping, requiring it to be excluded from the default measurement CPU list. As a result, you can now not only add, but also remove arbitrary CPUs from the default measurement CPU list in a flexible way. Removing takes precedence over adding. This rule applies to both, CPUs specified with +/- signs and to those defined with **--measurement-run-on-isolcpus**.

Jira:RHEL-9912[1]

rtla rebased to version 6.6 of the upstream kernel source code

The **rtla** utility has been upgraded to the latest upstream version, which provides multiple bug fixes and enhancements. Notable changes include:

- Added the **-C** option to specify additional control groups for **rtla** threads to run in, apart from the main **rtla** thread.

- Added the **--house-keeping** option to place **rtla** threads on a housekeeping CPU and to put measurement threads on different CPUs.

- Added support to the **timerlat** tracer so that you can run **timerlat hist** and **timerlat top** threads in user space.

Jira:RHEL-10079[1]

cyclicdeadline now supports generating a histogram of latencies

With this release, the **cyclicdeadline** utility supports generating a histogram of latencies. You can use this feature to get more insight into the frequency of latency spikes of different sizes, rather than getting just one worst-case number.

Jira:RHEL-9910[1]

SGX is now fully supported

Software Guard Extensions (SGX) is an Intel® technology for protecting software code and data from disclosure and modification.

The RHEL kernel provides the SGX version 1 and 2 functionality. Version 1 enables platforms using the **Flexible Launch Control** mechanism to use the SGX technology. Version 2 adds **Enclave Dynamic Memory Management** (EDMM). Notable features include:

- Modifying EPCM permissions of regular enclave pages that belong to an initialized enclave.

- Dynamic addition of regular enclave pages to an initialized enclave.

- Expanding an initialized enclave to accommodate more threads.

- Removing regular and TCS pages from an initialized enclave.

In this release, SGX moves from Technology Preview to a fully supported feature.

Bugzilla:2041883[1]

The Intel data streaming accelerator driver is now fully supported

The Intel data streaming accelerator driver (IDXD) is a kernel driver that provides an Intel CPU integrated accelerator. It includes a shared work queue with process address space ID (**pasid**) submission and shared virtual memory (SVM).

In this release, IDXD moves from a Technology Preview to a fully supported feature.

Jira:RHEL-10097[1]

The eBPF facility has been rebased to Linux kernel version 6.6

Notable changes and enhancements include:

- New dynamic pointers (**dynptrs**) of the **skb** and **xdp** type, which enable for more ergonomic and less brittle iteration through data and variable-sized accesses in BPF programs.

- A new BPF **netfilter** program type and minimal support to hook BPF programs to **netfilter** hooks, such as prerouting or forward.

- Multiple improvements to kernel pointers (**kptrs**):

  - You can use **kptrs** in more map types.

  - RCU semantics are enabled for task **kptrs**.

- New reference-counted local **kptrs** useful for adding a node to both the BPF **list** and **rbtree**.

- At load time, BPF programs can detect whether a particular **kfunc** exists or not.

- Several new **kfuncs** for working with **dynptrs**, **cgroups**, **sockets**, and **cpumasks**.

- New BPF links for attaching multiple **uprobes** and **usdt** probes, which is significantly faster and saves extra file descriptors (FDs).

- The BPF **map** element count is enabled for all program types.

- The memory usage reporting for all BPF **map** types is more precise.

- The **bpf_fib_lookup** BPF helper includes the routing table ID.

- The **BPF_OBJ_PIN** and **BPF_OBJ_GET** commands support **O_PATH** FDs.

Jira:RHEL-10691[1]

The libbpf-tools package is now available on IBM Z

The **libbpf-tools** package, which provides command line tools for the BPF Compiler Collection (BCC), is now available on the IBM Z architecture. As a result, you can now use commands from **libbpf-tools** on IBM Z.

Jira:RHEL-16325[1]

4.8. BOOT LOADER

DEP/NX support in the pre-boot stage

The memory protection feature known as Data Execution Prevention (DEP), No Execute (NX), or Execute Disable (XD), blocks the execution of code that is marked as non-executable. DEP/NX has been available in RHEL at the operating system level.

This release adds DEP/NX support in the GRUB and **shim** boot loaders. This can prevent certain vulnerabilities during the pre-boot stage, such as a malicious EFI driver that might execute certain attacks without the DEP/NX protection.

Jira:RHEL-10288[1]

4.9. FILE SYSTEMS AND STORAGE

Setting a filesystem size limit is now supported

With this update, users can now set a filesystem size limit when creating or modifying a filesystem. The **stratisd** service enables dynamic filesystem growth, but excessive expansion of an XFS filesystem can cause significant performance issues. The addition of this feature addresses potential performance issues that might occur when growing XFS filesystems beyond a certain threshold. By setting a filesystem size limit, users can prevent such issues and ensure optimal performance. Additionally, this feature enables better pool monitoring and maintenance by allowing users to impose an upper limit on a filesystem's size, ensuring efficient resource allocation.

Jira:RHEL-12898

Converting a standard LV to a thin LV by using lvconvert is now possible

By specifying a standard logical volume (LV) as a thin pool data, you can now convert a standard LV to a thin LV by using the **lvconvert** command. With this update, you can convert existing LVs to use the thin provisioning facility.

Jira:RHEL-8357

multipathd now supports detecting FPIN-Li events for NVMe devices

Previously, the **multipathd** command would only monitor Integrity Fabric Performance Impact Notification (PFIN-Li) events on SCSI devices. **multipathd** could listen for Link Integrity events sent by a Fibre Channel fabric and use it to mark paths as marginal. This feature was only supported for multipath devices on top of SCSI devices, and **multipathd** was unable to mark Non-volatile Memory Express (NVMe) device paths as marginal by limiting the use of this feature.

With this update, **multipathd** supports detecting FPIN-Li events for both SCSI and NVMe devices. As a result, multipath now does not use paths without a good fabric connection, while other paths are available. This helps to avoid IO delays in such situations.

Jira:RHEL-6678

max_retries option is now added to the defaults section of multipath.conf

This enhancement adds the **max_retries** option to the **defaults** section of the **multipath.conf** file. By default this option is unset, and uses the SCSI layer's default value of 5 retries. The valid values for this option is from **0** to **5**. When this option is set, it overrides the default value of the **max_retries sysfs** attribute for SCSI devices. This attribute controls the number of times the SCSI layer retries I/O commands before returning failure when it encounters certain error types.

If users encounter an issue where multipath's path checkers return success but I/O to a device is hanging, they can set this option to decrease the time before the I/O will be retried down another path.

Jira:RHEL-1729[1]

auto_resize option is now added to the defaults section of multipath.conf

Previously, to resize a multipath device, you had to manually execute the **multipathd resize map <name>** command. With this update, the **auto_resize** option is now added to the **defaults** section of the **multipath.conf** file. This option controls when the **multipathd** command can automatically resize a multipath device. The following are the different values for **auto_resize**:

- By default, **auto_resize** is set to **never**. In this case, **multipathd** works without any change.

- If **auto_resize** is set to **grow_only**, **multipathd** automatically resizes the multipath device when the device's paths have grown in size.

- If **auto_resize** is set to **grow_shrink**, **multipathd** automatically shrinks the multipath device when the device's paths are decreased in size.

As a result, when this option is enabled, you no longer need to manually resize your multipath devices.

Jira:RHEL-986[1]

Changes to Arcus NVMeoFC multipath.conf settings are now included in kernel

Device-mapper-multipath now has a built-in configuration for the HPE Alletra 9000 NVMeFC array. Arcus added support for ANA (Asymmetric Namespace Access) for NVMeoFC. This is similar to ALUA

for SCSI. A change in the **multipath.conf** is required for a RHEL host to use this feature and send only I/O to ANA optimized paths when available. Without this change, device mapper was sending I/O to both ANA optimized and ANA non-optimized paths.

> **NOTE**
>
> This change is only for NVMeoFC. FCP **multipath.conf** content already had this setting for supporting ALUA previously.

Jira:RHEL-1830

stratis-cli rebased to version 3.6.0

The **stratis-cli** package has been upgraded to version 3.6.0. Notable bug fixes and enhancements include:

- The **stratis-cli** command-line interface supports an additional option to set the file system size limit on creation. The **set-size-limit** and **unset-size-limit** are two new file system commands, which sets or unsets the file system size limit after creating a file system.

- **stratis-cli** now incorporates password verification when it is used to set a key in the kernel keyring by using a manual entry.

- **stratis-cli** now supports specifying a pool either by name or by UUID when stopping a pool.

- **stratis-cli** also gets updates with various internal improvements, and now enforces a requirement of at least the python 3.9 version in its package configuration.

Jira:RHEL-2265[1]

boom rebased to version 1.6.0

The **boom** package has been upgraded to version 3.6.0. Notable enhancements include:

- Support for multi-volume snapshot boot syntax supported by the **systemd** command.

- The **new --mount** and **--no-fstab** options are added to specify additional volumes to mount at the boot entry.

Jira:RHEL-16813

NVMe-FC Boot from SAN is now fully supported

The Non-volatile Memory Express (NVMe) over Fibre Channel (NVMe/FC) Boot, which was introduced in Red Hat Enterprise Linux 9.2 as a Technology Preview, is now fully supported. Some NVMe/FC host bus adapters support a NVMe/FC boot capability. For more information on programming a Host Bus Adapter (HBA) to enable NVMe/FC boot capability, see the NVMe/FC host bus adapter manufacturer's documentation.

Jira:RHEL-1492[1]

4.10. HIGH AVAILABILITY AND CLUSTERS

pcs support for ISO 8601 duration specification for time properties

The **pcs** command-line interface now allows you to specify values for Pacemaker time properties according to the ISO 8601 duration specification standard.

Jira:RHEL-7672

Support for new `pscd` Web UI features

The **pscd** Web UI now supports the following features:

- Moving a cluster resource off the node on which it is currently running

- Banning a resource from running on a node

- Displaying cluster status that shows the age of the cluster status and when the cluster state is being reloaded

- Requesting a reload of the cluster status display

Jira:RHEL-7582, Jira:RHEL-7739

TLS cipher list now defaults to system-wide crypto policy

Previously, the **pcsd** TLS cipher list was set to **DEFAULT:!RC4:!3DES:@STRENGTH** by default. With this update, the cipher list is defined by the system-wide crypto policy by default. The TLS ciphers accepted by the **pcsd** daemon might change with this upgrade, depending on the currently set crypto policy. For more information about the crypto policies framework, see the **crypto-policies**(7) man page.

Jira:RHEL-7724

4.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Python 3.12 available in RHEL 9

RHEL 9.4 introduces Python 3.12, provided by the new package **python3.12** and a suite of packages built for it, as well as the **ubi9/python-312** container image.

Notable enhancements compared to the previously released Python 3.11 include:

- Python introduces a new **type** statement and new type parameter syntax for generic classes and functions.

- Formatted string literal (f-strings) have been formalized in the grammar and can now be integrated into the parser directly.

- Python now provides a unique per-interpreter global interpreter lock (GIL).

- You can now use the buffer protocol from Python code.

- To improve security, the builtin **hashlib** implementations of the SHA1, SHA3, SHA2-384, SHA2-512, and MD5 cryptographic algorithms have been replaced with formally verified code from the HACL* project. The builtin implementations remain available as fallback if OpenSSL does not provide them.

- Dictionary, list, and set comprehensions in **CPython** are now inlined. This significantly increases the speed of a comprehension execution.

- **CPython** now supports the Linux **perf** profiler.

- **CPython** now provides stack overflow protection on supported platforms.

Python 3.12 and packages built for it can be installed in parallel with Python 3.9 and Python 3.11 on the same system.

To install packages from the **python3.12** stack, use, for example:

```
# dnf install python3.12
# dnf install python3.12-pip
```

To run the interpreter, use, for example:

```
$ python3.12
$ python3.12 -m pip --help
```

See Installing and using Python for more information.

For information about the length of support of Python 3.12, see Red Hat Enterprise Linux Application Streams Life Cycle.

Jira:RHEL-14941

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

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

This update introduces a new **PYTHON_EMAIL_DISABLE_STRICT_ADDR_PARSING** environment variable. When you set this variable to **true**, the previous, less strict parsing behavior is the default for the entire system:

```
export PYTHON_EMAIL_DISABLE_STRICT_ADDR_PARSING=true
```

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

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

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

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

Jira:RHELDOCS-17369[1]

A new module stream: ruby:3.3

RHEL 9.4 introduces Ruby 3.3.0 in a new **ruby:3.3** module stream. This version provides a number of performance improvements, bug and security fixes, and new features over **Ruby 3.1** distributed with RHEL 9.1.

Notable enhancements include:

- You can use the new **Prism** parser instead of **Ripper**. **Prism** is a portable, error tolerant, and maintainable recursive descent parser for the Ruby language.

- YJIT, the Ruby just-in-time (JIT) compiler implementation, is no longer experimental and it provides major performance improvements.

- The **Regexp** matching algorithm has been improved to reduce the impact of potential Regular Expression Denial of Service (ReDoS) vulnerabilities.

- The new experimental RJIT (a pure-Ruby JIT) compiler replaces MJIT. Use YJIT in production.

- A new M:N thread scheduler is now available.

Other notable changes:

- You must now use the **Lrama** LALR parser generator instead of **Bison**.

- Several deprecated methods and constants have been removed.

- The **Racc** gem has been promoted from a default gem to a bundled gem.

To install the **ruby:3.3** module stream, use:

```
# dnf module install ruby:3.3
```

If you want to upgrade from an earlier **ruby** module stream, see Switching to a later stream .

For information about the length of support of Ruby 3.3, see Red Hat Enterprise Linux Application Streams Life Cycle.

Jira:RHEL-17089[1]

A new module stream: php:8.2

RHEL 9.4 adds PHP 8.2 as a new **php:8.2** module stream.

Improvements in this release include:

- Readonly classes

- Several new stand-alone types

- A new **Random** extension

- Constraints in traits

To install the **php:8.2** module stream, use the following command:

```
# dnf module install php:8.2
```

If you want to upgrade from the **php:8.1** stream, see Switching to a later stream .

For details regarding PHP usage on RHEL 9, see Using the PHP scripting language .

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

Jira:RHEL-14699[1]

A new module stream: nginx:1.24

The nginx 1.24 web and proxy server is now available as the **nginx:1.24** module stream. This update provides a number of bug fixes, security fixes, new features, and enhancements over the previously released version 1.22.

New features and changes related to Transport Layer Security (TLS):

- Encryption keys are now automatically rotated for TLS session tickets when using shared memory in the **ssl_session_cache** directive.

- Memory usage has been optimized in configurations with Secure Sockets Layer (SSL) proxy.

- You can now disable looking up IPv4 addresses while resolving by using the **ipv4=off** parameter of the **resolver** directive.

- nginx now supports the **$proxy_protocol_tlv_*** variables, which store the values of the Type-Length-Value (TLV) fields that appear in the PROXY v2 TLV protocol.

- The **ngx_http_gzip_static_module** module now supports byte ranges.

Other changes:

- Header lines are now represented as linked lists in the internal API.

- nginx now concatenates identically named header strings passed to the FastCGI, SCGI, and uwsgi back ends in the **$r->header_in()** method of the **ngx_http_perl_module**, and during lookups of the **$http_...**, **$sent_http_...**, **$sent_trailer_...**, **$upstream_http_...**, and **$upstream_trailer_...** variables.

- nginx now displays a warning if protocol parameters of a listening socket are redefined.

- nginx now closes connections with lingering if pipelining was used by the client.

- The logging level of various SSL errors has been lowered, for example, from **Critical** to **Informational**.

To install the **nginx:1.24** stream, use:

```
# dnf module install nginx:1.24
```

To upgrade from the **nginx 1.22** stream, switch to a later stream .

For more information, see Setting up and configuring NGINX.

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

Jira:RHEL-14713[1]

A new module stream: mariadb:10.11

MariaDB 10.11 is now available as a new module stream, **mariadb:10.11**. Notable enhancements over the previously available version 10.5 include:

- A new **sys_schema** feature.

- Atomic Data Definition Language (DDL) statements.

- A new **GRANT ... TO PUBLIC** privilege.

- Separate **SUPER** and **READ ONLY ADMIN** privileges.

- A new **UUID** database data type.

- Support for the Secure Socket Layer (SSL) protocol version 3; the MariaDB server now requires correctly configured SSL to start.

- Support for the natural sort order through the **natural_sort_key()** function.

- A new **SFORMAT** function for arbitrary text formatting.

- Changes to the UTF-8 charset and the UCA-14 collation.

- **systemd** socket activation files available in the **/usr/share/** directory. Note that they are not a part of the default configuration in RHEL as opposed to upstream.

- Error messages containing the **MariaDB** string instead of **MySQL**.

- Error messages available in the Chinese language.

- Changes to the default logrotate file.

- For MariaDB and MySQL clients, the connection property specified on the command line (for example, **--port=3306**), now forces the protocol type of communication between the client and the server, such as **tcp**, **socket**, **pipe**, or **memory**.

For more information about changes in MariaDB 10.11, see Notable differences between MariaDB 10.5 and MariaDB 10.11.

For more information about MariaDB, see Using MariaDB.

To install the **mariadb:10.11** stream, use:

```
# dnf module install mariadb:10.11
```

If you want to upgrade from MariaDB 10.5, see Upgrading from MariaDB 10.5 to MariaDB 10.11 .

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

Jira:RHEL-3638

A new module stream: **postgresql:16**

RHEL 9.4 introduces PostgreSQL 16 as the **postgresql:16** module stream. PostgreSQL 16 provides a number of new features and enhancements over version 15.

Notable enhancements include:

- Enhanced bulk loading improves performance.

- The **libpq** library now supports connection-level load balancing. You can use the new **load_balance_hosts** option for more efficient load balancing.

- You can now create custom configuration files and include them in the **pg_hba.conf** and **pg_ident.conf** files.

- PostgreSQL now supports regular expression matching on database and role entries in the **pg_hba.conf** file.

Other changes include:

- PostgreSQL is no longer distributed with the **postmaster** binary. Users who start the **postgresql** server by using the provided **systemd** unit file (the **systemctl start postgres** command) are not affected by this change. If you previously started the **postgresql** server directly through the **postmaster** binary, you must now use the **postgres** binary instead.

- PostgreSQL no longer provides documentation in PDF format within the package. Use the online documentation instead.

See also Using PostgreSQL.

To install the **postgresql:16** stream, use the following command:

```
# dnf module install postgresql:16
```

If you want to upgrade from an earlier **postgresql** stream within RHEL 9, follow the procedure described in Switching to a later stream and then migrate your PostgreSQL data as described in Migrating to a RHEL 9 version of PostgreSQL.

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

Jira:RHEL-3635

Git rebased to version 2.43.0

The Git version control system has been updated to version 2.43.0, which provides bug fixes, enhancements, and performance improvements over the previously released version 2.39.

Notable enhancements include:

- You can now use the new **--source** option with the **git check-attr** command to read the **.gitattributes** file from the provided tree-ish object instead of the current working directory.

- Git can now pass information from the **WWW-Authenticate** response-type header to credential helpers.

- In case of an empty commit, the **git format-patch** command now writes an output file containing a header of the commit instead of creating an empty file.

- You can now use the **git blame --contents=*<file> <revision> -- <path>*** command to find the origins of lines starting at *<file>* contents through the history that leads to *<revision>*.

- The **git log --format** command now accepts the **%(decorate)** placeholder for further customization to extend the capabilities provided by the **--decorate** option.

Jira:RHEL-17100[1]

Git LFS rebased to version 3.4.1

The Git Large File Storage (LFS) extension has been updated to version 3.4.1, which provides bug fixes, enhancements, and performance improvements over the previously released version 3.2.0.

Notable changes include:

- The **git lfs push** command can now read references and object IDs from standard input.

- Git LFS now handles alternative remotes without relying on Git.

- Git LFS now supports the **WWW-Authenticate** response-type header as a credential helper.

Jira:RHEL-17101[1]

4.12. COMPILERS AND DEVELOPMENT TOOLS

LLVM Toolset rebased to version 17.0.6

LLVM Toolset has been updated to version 17.0.6.

Notable enhancements include:

- The opaque pointers migration is now completed.

- Removed support for the legacy pass manager in middle-end optimization.

Clang changes:

- C++20 coroutines are no longer considered experimental.

- Improved code generation for the **std::move** function and similar in unoptimized builds.

For more information, see the LLVM and Clang upstream release notes.

Jira:RHEL-9283

Rust Toolset rebased to version 1.75.0

Rust Toolset has been updated to version 1.75.0.

Notable enhancements include:

- Constant evaluation time is now unlimited

- Cleaner panic messages

- Cargo registry authentication

- **async fn** and opaque return types in traits

Jira:RHEL-12963

Go Toolset rebased to version 1.21.0

Go Toolset has been updated to version 1.21.0.

Notable enhancements include:

- **min**, **max**, and **clear** built-ins have been added.

- Official support for profile guided optimization has been added.

- Package initialization order is now more precisely defined.

- Type inferencing is improved.

- Backwards compatibility support is improved.

For more information, see the Go upstream release notes.

Jira:RHEL-11871[1]

Clang resource directory moved

The Clang resource directory, where Clang stores its internal headers and libraries, has been moved from **/usr/lib64/clang/17** to **/usr/lib/clang/17**.

Jira:RHEL-9346

elfutils rebased to version 0.190

The **elfutils** package has been updated to version 0.190. Notable improvements include:

- The **libelf** library now supports relative relocation (RELR).

- The **libdw** library now recognizes **.debug_[ct]u_index** sections.

- The **eu-readelf** utility now supports a new **-Ds**, **--use-dynamic --symbol** option to show symbols through the dynamic segment without using ELF sections.

- The **eu-readelf** utility can now show **.gdb_index** version 9.

- A new **eu-scrlines** utility compiles a list of source files associated with a specified DWARF or ELF file.

- A **debuginfod** server schema has changed for a 60% compression in file name representation (this requires reindexing).

Jira:RHEL-12489

systemtap rebased to version 5.0

The **systemtap** package has been updated to version 5.0. Notable enhancements include:

- Faster and more reliable kernel-user transport.

- Extended DWARF5 debuginfo format support.

Jira:RHEL-12488

Updated GCC Toolset 13

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

Notable changes introduced in RHEL 9.4 include:

- The GCC compiler has been updated to version 13.2.1, which provides many bug fixes and enhancements that are available in upstream GCC.

- GCC and **binutils** now support AMD CPUs based on the **znver5** core through the **-march=znver5** compiler switch.

- **annobin** has been updated to version 12.32.

- The **annobin** plugin for GCC now defaults to using a more compressed format for the notes that it stores in object files, resulting in smaller object files and faster link times, especially in large, complex programs.

The following tools and versions are provided by GCC Toolset 13: GCC:: 13.2.1 GDB:: 12.1 binutils:: 2.40 dwz:: 0.14 annobin:: 12.32

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

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

To run a tool from GCC Toolset 13:

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

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

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

For more information, see GCC Toolset 13.

Jira:RHEL-23798[1]

Compiling with GCC and the -fstack-protector flag no longer fails to guard dynamic stack allocations on 64-bit ARM

Previously, on the 64-bit ARM architecture, the system GCC compiler with the **-fstack-protector** flag failed to detect a buffer overflow in functions containing a C99 variable-length array or an **alloca()**-allocated object. Consequently, an attacker could overwrite saved registers on the stack. With this update, the buffer overflow detection on 64-bit ARM has been fixed. As a result, applications compiled with the system GCC are more secure.

Jira:RHEL-17638[1]

GCC Toolset 13: Compiling with GCC and the -fstack-protector flag no longer fails to guard dynamic stack allocations on 64-bit ARM

Previously, on the 64-bit ARM architecture, the GCC compiler with the **-fstack-protector** flag failed to detect a buffer overflow in functions containing a C99 variable-length array or an **alloca()**-allocated object. Consequently, an attacker could overwrite saved registers on the stack. With this update, the buffer overflow detection on 64-bit ARM has been fixed. As a result, applications compiled with GCC are more secure.

Jira:RHEL-16998

pcp updated to version 6.2.0

The **pcp** package has been updated to version 6.2.0. Notable improvements include:

- **pcp-htop** now supports user-defined tabs.

- **pcp-atop** now supports a new bar graph visualization mode.

- OpenMetrics PMDA metric labels and logging are improved.

- Additional Linux kernel virtual memory metrics have been added.

- New tools:

  - **pmlogredact**

  - **pcp-buddyinfo**

  - **pcp-meminfo**

  - **pcp-netstat**

  - **pcp-slabinfo**

  - **pcp-zoneinfo**

Jira:RHEL-2317[1]

A new **grafana-selinux** package

Previously, the default installation of **grafana-server** ran as an **unconfined_service_t** SELinux type. This update adds the new **grafana-selinux** package, which contains an SELinux policy for **grafana-server** and which is installed by default with **grafana-server**. As a result, **grafana-server** now runs as **grafana_t** SELinux type.

Jira:RHEL-7505

papi supports new processor microarchitectures

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

- AMD Zen 4

- 4th Generation Intel® Xeon® Scalable Processors

Jira:RHEL-9333[1], Jira:RHEL-9335, Jira:RHEL-9334

New package: **maven-openjdk21**

The **maven:3.8** module stream now includes the **maven-openjdk21** subpackage, which provides the Maven JDK binding for OpenJDK 21 and configures Maven to use the system OpenJDK 21.

Jira:RHEL-13046[1]

New package: **libzip-tools**

RHEL 9.4 introduces the **libzip-tools** package, which provides utilities such such as **zipcmp**, **zipmerge**, and **ziptool**.

Jira:RHEL-17567

cmake rebased to version 3.26

The **cmake** package has been updated to version 3.26. Notable improvements include:

- Added support for the C17 and C18 language standards.

- **cmake** can now query the **/etc/os-release** file for operating system identification information.

- Added support for the CUDA 20 and **nvtx3** libraries.

- Added support for the Python stable application binary interface.

- Added support for Perl 5 in the Simplified Wrapper and Interface Generator (SWIG) tool.

Jira:RHEL-7393

4.13. IDENTITY MANAGEMENT

A new passwordless authentication method is available in SSSD

With this update, you can enable and configure passwordless authentication in SSSD to use a biometric device that is compatible with the FIDO2 specification, for example a YubiKey. You must register the FIDO2 token in advance and store this registration information in the user account in RHEL IdM, Active Directory, or an LDAP store. RHEL implements FIDO2 compatibility with the **libfido2** library, which currently only supports USB-based tokens.

Jira:RHELDOCS-17841[1]

The **ansible-freeipa ipauser** and **ipagroup** modules now support a new **renamed** state

With this update, you can use the **renamed** state in **ansible-freeipa ipauser** module to change the user name of an existing IdM user. You can also use this state in **ansible-freeipa ipagroup** module to change the group name of an existing IdM group.

Jira:RHEL-4962

Identity Management users can now use external identity providers to authenticate to IdM

With this enhancement, you can now associate Identity Management (IdM) users with external identity providers (IdPs) that support the OAuth 2 device authorization flow. Examples of such IdPs include Red Hat build of Keycloak, Azure Entra ID, Github, Google, and Facebook.

If an IdP reference and an associated IdP user ID exist in IdM, you can use them to enable an IdM user to authenticate at the external IdP. After performing authentication and authorization at the external IdP, the IdM user receives a Kerberos ticket with single sign-on capabilities. The user must authenticate with the SSSD version available in RHEL 9.1 or later.

Jira:RHELPLAN-169666[1]

ipa rebased to version 4.11

The **ipa** package has been updated from version 4.10 to 4.11. Notable changes include:

- Support for FIDO2-based passkeys.

- Initial implementation of resource-based constrained delegation (RBCD) for Kerberos services.

- Context manager for **ipalib.api** to automatically configure, connect, and disconnect.

- The installation of an IdM replica now occurs against a chosen server, not only for Kerberos authentication but also for all IPA API and CA requests.

- The **ansible-freeipa** package has been rebased from version 1.11 to 1.12.1.

- The **ipa-healthcheck** package has been rebased from version 0.12 to 0.16.

For more information, see the upstream release notes.

Jira:RHEL-11652

Deleting expired KCM Kerberos tickets

Previously, if you attempted to add a new credential to the Kerberos Credential Manager (KCM) and you had already reached the storage space limit, the new credential was rejected. The user storage space is limited by the **max_uid_ccaches** configuration option that has a default value of 64. With this update, if you have already reached the storage space limit, your oldest expired credential is removed and the new credential is added to the KCM. If there are no expired credentials, the operation fails and an error is returned. To prevent this issue, you can free some space by removing credentials using the **kdestroy** command.

Jira:SSSD-6216

IdM now supports the idoverrideuser, idoverridegroup and idview Ansible modules

With this update, the **ansible-freeipa** package now contains the following modules:

idoverrideuser

Allows you to override user attributes for users stored in the Identity Management (IdM) LDAP server, for example, the user login name, home directory, certificate, or SSH keys.

idoverridegroup

Allows you to override attributes for groups stored in the IdM LDAP server, for example, the name of the group, its GID, or description.

idview

Allows you to organize user and group ID overrides and apply them to specific IdM hosts.

In the future, you will be able to use these modules to enable AD users to use smart cards to log in to IdM.

Jira:RHEL-16934

The idp Ansible module allows associating IdM users with external IdPs

With this update, you can use the **idp ansible-freeipa** module to associate Identity Management (IdM) users with external identity providers (IdP) that support the OAuth 2 device authorization flow. If an IdP reference and an associated IdP user ID exist in IdM, you can use them to enable IdP authentication for an IdM user.

After performing authentication and authorization at the external IdP, the IdM user receives a Kerberos ticket with single sign-on capabilities. The user must authenticate with the SSSD version available in RHEL 8.7 or later.

Jira:RHEL-16939

getcert add-ca returns a new return code if a certificate is already present or tracked

With this update, the **getcert** command returns a specific return code, **2**, if you try to add or track a certificate that is already present or tracked. Previously, the command returned return code **1** on any error condition.

Jira:RHEL-22302

The delegation of DNS zone management is now enabled in ansible-freeipa

You can now use the **dnszone ansible-freeipa** module to delegate DNS zone management. Use the **permission** or **managedby** variable of the **dnszone** module to configure a per-zone access delegation permission.

Jira:RHEL-19134

Enforcing OTP usage for all LDAP clients

With the release of the RHBA-2024:2558 advisory, in RHEL IdM, you can now set the default behavior for LDAP server authentication of user accounts with two-factor (OTP) authentication configured. If OTP is enforced, LDAP clients cannot authenticate against an LDAP server using single factor authentication (a password) for users that have associated OTP tokens. This method is already enforced through the Kerberos backend by using a special LDAP control with OID 2.16.840.1.113730.3.8.10.7 without any data.

- To enforce OTP usage for all LDAP clients, administrators can use the following command:

  ```
  $ ipa config-mod --addattr ipaconfigstring=EnforceLDAPOTP
  ```

- To change back to the previous OTP behavior for all LDAP clients, use the following command:

  ```
  $ ipa config-mod --delattr ipaconfigstring=EnforceLDAPOTP
  ```

Jira:RHEL-23377[1]

The runasuser_group parameter is now available in ansible-freeipa ipasudorule

With this update, you can set Groups of RunAs Users for a **sudo** rule by using the **ansible-freeipa ipasudorule** module. The option is already available in the Identity Management (IdM) command-line interface and the IdM Web UI.

Jira:RHEL-19130

389-ds-base rebased to version 2.4.5

The **389-ds-base** package has been updated to version 2.4.5. Notable bug fixes and enhancements over version 2.3.4 include:

- https://www.port389.org/docs/389ds/releases/release-2-3-5.html

- https://www.port389.org/docs/389ds/releases/release-2-3-6.html

- https://www.port389.org/docs/389ds/releases/release-2-3-7.html

- https://www.port389.org/docs/389ds/releases/release-2-4-0.html

- https://www.port389.org/docs/389ds/releases/release-2-4-1.html

- https://www.port389.org/docs/389ds/releases/release-2-4-2.html

- https://www.port389.org/docs/389ds/releases/release-2-4-3.html

- https://www.port389.org/docs/389ds/releases/release-2-4-4.html

- https://www.port389.org/docs/389ds/releases/release-2-4-5.html

Jira:RHEL-15907

Transparent Huge Pages are now disabled by default for the ns-slapd process

When large database caches are used, Transparent Huge Pages (THP) can have a negative effect on Directory Server performance under heavy load, for example, high memory footprint, high CPU usage and latency spikes. With this enhancement, a new **THP_DISABLE=1** configuration option was added to the **/usr/lib/systemd/system/dirsrv@.service.d/custom.conf** drop-in configuration file for the **dirsrv systemd** unit to disable THP for the **ns-slapd** process.

In addition, the Directory Server health check tool now detects the THP settings. If you enabled THP system-wide and for the Directory Server instance, the health check tool informs you about the enabled THP and prints recommendations on how to disable them.

Jira:RHEL-5142

The new lastLoginHistSize configuration attribute is now available for the Account Policy plug-in

Previously, when a user did a successful bind, only the time of the last login was available. With this update, you can use the new **lastLoginHistSize** configuration attribute to manage a history of successful logins. By default, the last five successful logins are saved.

Note that for the **lastLoginHistSize** attribute to collect statistics of successful logins, you must enable the **alwaysRecordLogin** attribute for the Account Policy plug-in.

For more details, see lastLoginHistSize.

Jira:RHEL-5133[1]

The new notes=M message in the access log to identify MFA binds

With this update, when you configure the two-factor authentication for user accounts by using a pre-bind authentication plug-in, such as MFA plug-in, the Directory Server log files record the following messages during **BIND** operations:

- The access log records the new **notes=M** note message:

  ```
  [time_stamp] conn=1 op=0 BIND dn="uid=jdoe,ou=people,dc=example,dc=com" method=128
  version=3
  [time_stamp] conn=1 op=0 RESULT err=0 tag=97 nentries=0 wtime=0.000111632
  optime=0.006612223 etime=0.006722325 notes=M details="Multi-factor Authentication"
  dn="uid=jdoe,ou=people,dc=example,dc=com"
  ```

- The security log records the new **SIMPLE/MFA** bind method:

  ```
  { "date": "[time_stamp] ", "utc_time": "1709327649.232748932", "event": "BIND_SUCCESS",
  "dn": "uid=djoe,ou=people,dc=example,dc=com", "bind_method": "SIMPLE\/MFA",
  "root_dn": false, "client_ip": "::1", "server_ip": "::1", "ldap_version": 3, "conn_id": 1, "op_id": 0,
  "msg": "" }
  ```

Note that for the access and security logs to record such messages, the pre-bind authentication plug-in must set the flag by using the SLAPI API if a bind was part of this plug-in.

Jira:RHELDOCS-17838[1]

The new inchainMatch matching rule is now available

With this update, a client application can use the new **inchainMatch** matching rule to search for the ancestry of an LDAP entry. The **member**, **manager**, **parentOrganization**, and **memberof** attributes can be used with the **inchainMatch** matching rule and the following searches can be performed:

- Find all direct or indirect groups in which a user is a member.

- Find all direct or indirect users whose manager is a certain user.

- Find all direct or indirect organizations an entry belongs to.

- Finds all direct or indirect members of a certain group.

Note that for performance reasons, you must index the **member**, **manager**, **parentOrganization**, and **memberof** attributes if the client application performs searches against these attributes by using the **inchainMatch** matching rule.

Directory Server uses the In Chain plug-in that is enabled by default to implement the **inchainMatch** matching rule. However, because **inchainMatch** is expensive to compute, an access control instruction (ACI) limits the matching rule usage.

For more details, refer to Using **inchainMatch** matching rule to find the ancestry of an LDAP entry .

Jira:RHELDOCS-17256[1]

The HAProxy protocol is now supported for the 389-ds-base package

Previously, Directory Server did not differentiate incoming connections between proxy and non-proxy clients. With this update, you can use the new **nsslapd-haproxy-trusted-ip** multi-valued configuration attribute to configure the list of trusted proxy servers. When **nsslapd-haproxy-trusted-ip** is configured under the **cn=config** entry, Directory Server uses the HAProxy protocol to receive client IP addresses via an additional TCP header so that access control instructions (ACIs) can be correctly evaluated and client traffic can be logged.

If an untrusted proxy server initiates a bind request, Directory Server rejects the request and records the following message to the error log file:

> [time_stamp] conn=5 op=-1 fd=64 Disconnect - Protocol error - Unknown Proxy - P4

For more details, see nsslapd-haproxy-trusted-ip.

Jira:RHEL-5130

samba rebased to version 4.19.4

The **samba** packages have been upgraded to upstream version 4.19.4, which provides bug fixes and enhancements over the previous version. The most notable changes are:

- Command-line options in the **smbget** utility have been renamed and removed for a consistent user experience. However, this can break existing scripts or jobs that use the utility. See the **smbget --help** command and **smbget(1)** man page for further details about the new options.

- If the **winbind debug traceid** option is enabled, the **winbind** service now logs, additionally, the following fields:

  - **traceid**: Tracks the records belonging to the same request.

  - **depth**: Tracks the request nesting level.

- Samba no longer uses its own cryptography implementations and, instead, now fully uses cryptographic functionality provided by the GnuTLS library.

- The **directory name cache size** option was removed.

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

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

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

Jira:RHEL-16476

Identity Management API is now fully supported

The Identity Management (IdM) API was available as a Technology Preview in RHEL 9.2. Since RHEL 9.3, it has been fully supported.

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

- Administrators can use previous or later versions of IdM on the server than on the managing client.

- Developers can use a specific version of an IdM call, even if the IdM version changes on the server.

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

NOTE

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

Bugzilla:1513934

4.14. THE WEB CONSOLE

RHEL web console can now generate Ansible and shell scripts

In the web console, you can now easily access and copy automation scripts on the **kdump** configuration page. You can then use the generated script to implement a specific **kdump** configuration on multiple systems.

Jira:RHELDOCS-17060[1]

Simplified managing storage and resizing partitions on Storage

The Storage section of the web console is now redesigned. The new design improved visibility across all views. The overview page now presents all storage objects in a comprehensive table, which makes it easier to perform operations directly. You can click any row to view detailed information and any supplementary actions. Additionally, you can now resize partitions from the Storage section.

Jira:RHELDOCS-17056[1]

4.15. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The ad_integration RHEL System Role now supports configuring dynamic DNS update options

With this update, the **ad_integration** RHEL System Role supports configuring options for dynamic DNS updates using SSSD when integrated with Active Directory (AD). By default, SSSD will attempt to automatically refresh the DNS record:

- When the identity provider comes online (always).

- At a specified interval (optional configuration); by default, the AD provider updates the DNS record every 24 hours.

You can change these and other settings using the new variables in the **ad_integration** System Role. For example, you can set **ad_dyndns_refresh_interval** to **172800** to change the DNS record refresh interval to 48 hours.

Jira:RHELDOCS-17372[1]

The Storage RHEL System Roles now support shared LVM device management

The RHEL System Roles now support the creation and management of shared logical volumes and volume groups.

Jira:RHEL-1535

Microsoft SQL Server 2022 available on RHEL 9

The **mssql-server** system role is now available on RHEL 9. The role adds two variables:

1. **mssql_run_selinux_confined** to control whether to run SQL Server as a confined application or not. If set to **true**, the role installs the **mssql-server-selinux** package. If set to **false**, the role removes the **mssql-server-selinux** package. Default setting is **true** for RHEL 9 managed nodes and **false** for other managed nodes.

2. **mssql_manage_selinux** to control whether to configure SELinux. When set to **true**, the variable configures the enforcing or permissive mode based on the value of the **mssql_run_selinux_confined** variable.

Jira:RHEL-16342

The rhc system role now supports RHEL 7 systems

You can now manage RHEL 7 systems by using the **rhc** system role. Register the RHEL 7 system to Red Hat Subscription Management (RHSM) and Insights and start managing your system using the **rhc** system role.

Using the **rhc_insights.remediation** parameter has no impact on RHEL 7 systems as the Insights Remediation feature is currently not available on RHEL 7.

Jira:RHEL-16976

New RHEL System Role for configuring **fapolicyd**

With the new **fapolicyd** RHEL System Role, you can use Ansible playbooks to manage and configure the **fapolicyd** framework. The **fapolicyd** software framework controls the execution of applications based on a user-defined policy.

Jira:RHEL-16542

The RHEL System Roles now support LVM snapshot management

With this enhancement, you can use the new **snapshot** RHEL System Role to create, configure, and manage LVM snapshots.

Jira:RHEL-16552

The Nmstate API and the **network** RHEL System role now support new route types

With this enhancement, you can use the following route types with the Nmstate API and the **network** RHEL System Role:

- **blackhole**

- **prohibit**

- **unreachable**

Jira:RHEL-19579[1]

The **ad_integration** RHEL System Role now supports custom SSSD domain configuration settings

Previously, when using the **ad_integration** RHEL System Role, it was not possible to add custom settings to the domain configuration section in the **sssd.conf** file using the role. With this enhancement, the **ad_integration** role can now modify the **sssd.conf** file and, as a result, you can use custom SSSD settings.

Jira:RHEL-17668

The **ad_integration** RHEL System Role now supports custom SSSD settings

Previously, when using the **ad_integration** RHEL System Role, it was not possible to add custom settings to the **[sssd]** section in the **sssd.conf** file using the role. With this enhancement, the **ad_integration** role can now modify the **sssd.conf** file and, as a result, you can use custom SSSD settings.

Jira:RHEL-21133

New rhc_insights.display_name option in the **rhc** role to set display names

You can now configure or update the display name of the system registered to Red Hat Insights by using the new **rhc_insights.display_name** parameter. The parameter allows you to name the system based on your preference to easily manage systems in the Insights Inventory. If your system is already

connected with Red Hat Insights, use the parameter to update the existing display name. If the display name is not set explicitly on registration, it is set to the hostname by default. It is not possible to automatically revert the display name to the hostname, but it can be set so manually.

Jira:RHEL-16964

New RHEL System Role for configuring **fapolicyd**

With the new **fapolicyd** RHEL System Role, you can use Ansible playbooks to manage and configure the **fapolicyd** framework. The **fapolicyd** software framework controls the execution of applications based on a user-defined policy.

Jira:RHEL-16541

New **logging_preserve_fqdn** variable for the **logging** RHEL System Role

Previously, it was not possible to configure a fully qualified domain name (FQDN) using the **logging** system role. This update adds the optional **logging_preserve_fqdn** variable, which you can use to set the **preserveFQDN** configuration option in **rsyslog** to use the full FQDN instead of a short name in syslog entries.

Jira:RHEL-15932

The **logging** role supports general queue and general action parameters in output modules

Previously, it was not possible to configure general queue parameters and general action parameters with the **logging** role. With this update, the **logging** RHEL System Role supports configuration of general queue parameters and general action parameters in output modules.

Jira:RHEL-15439

The **postgresql** RHEL System Role now supports PostgreSQL 16

The **postgresql** RHEL System Role, which installs, configures, manages, and starts the PostgreSQL server, now supports PostgreSQL 16.

For more information about this system role, see Installing and configuring PostgreSQL by using the postgresql RHEL System Role.

Jira:RHEL-18962

Support for creation of volumes without creating a file system

With this enhancement, you can now create a new volume without creating a file system by specifying the **fs_type=unformatted** option.

Similarly, existing file systems can be removed using the same approach by ensuring that the safe mode is disabled.

Jira:RHEL-16212

Support for new **ha_cluster** System Role features

The **ha_cluster** System Role now supports the following features:

- Enablement of the repositories containing resilient storage packages, such as **dlm** or **gfs2**. A Resilient Storage subscription is needed to access the repository.

- Configuration of fencing levels, allowing a cluster to use multiple devices to fence nodes.

- Configuration of node attributes.

For information about the parameters you configure to implement these features, see Configuring a high-availability cluster by using the ha_cluster RHEL System Role.

Jira:RHEL-15876[1], Jira:RHEL-22106, Jira:RHEL-15910

ForwardToSyslog flag is now supported in the journald system role

In the **journald** RHEL System Role, the **journald_forward_to_syslog** variable controls whether the received messages should be forwarded to the traditional **syslog** daemon or not. The default value of this variable is **false**. With this enhancement, you can now configure the **ForwardToSyslog** flag by setting **journald_forward_to_syslog** to **true** in the inventory. As a result, when using remote logging systems such as Splunk, the logs are available in the **/var/log** files.

Jira:RHEL-21117

New rhc_insights.ansible_host option in the rhc role to set Ansible hostnames

You can now configure or update the Ansible hostname for the systems registered to Red Hat Insights by using the new **rhc_insights.ansible_host** parameter. When set, the parameter changes the **ansible_host** configuration in the **/etc/insights-client/insights-client.conf** file to your selected Ansible hostname. If your system is already connected with Red Hat Insights, this parameter will update the existing Ansible hostname.

Jira:RHEL-16974

New mssql_ha_prep_for_pacemaker variable

Previously, the **microsoft.sql.server** RHEL System Role did not have a variable to control whether to configure SQL Server for Pacemaker. This update adds the **mssql_ha_prep_for_pacemaker**. Set the variable to **false** if you do not want to configure your system for Pacemaker and you want to use another HA solution.

Jira:RHEL-19091

The sshd role now configures certificate-based SSH authentications

With the **sshd** RHEL System Role, you can now configure and manage multiple SSH servers to authenticate by using SSH certificates. This makes SSH authentications more secure because certificates are signed by a trusted CA and provide fine-grained access control, expiration dates, and centralized management.

Jira:RHEL-5972

Use the logging_max_message_size parameter instead of rsyslog_max_message_size in the logging system role

Previously, even though the **rsyslog_max_message_size** parameter was not supported, the **logging** RHEL System Role was using **rsyslog_max_message_size** instead of using the **logging_max_message_size** parameter. This enhancement ensures that **logging_max_message_size** is used and not **rsyslog_max_message_size** to set the maximum size for the log messages.

Jira:RHEL-15037

ratelimit_burst variable is only used if ratelimit_interval is set in logging system role

Previously, in the **logging** RHEL System Role, when the **ratelimit_interval** variable was not set, the role would use the **ratelimit_burst** variable to set the rsyslog **ratelimit.burst** setting. But it had no effect because it is also required to set **ratelimit_interval**.

With this enhancement, if **ratelimit_interval** is not set, the role does not set **ratelimit.burst**. If you want to set **ratelimit.burst**, you must set both **ratelimit_interval** and **ratelimit_burst** variables.

Jira:RHEL-19046

selinux role now prints a message when specifying a non-existent module

With this release, the **selinux** RHEL System Role prints an error message when you specify a non-existent module in the **selinux_modules.path** variable.

Jira:RHEL-19043

selinux role now supports configuring SELinux in disabled mode

With this update, the **selinux** RHEL System Role supports configuring SELinux ports, file contexts, and boolean mappings on nodes that have SELinux set to disabled. This is useful for configuration scenarios before you enable SELinux to permissive or enforcing mode on a system.

Jira:RHEL-15870

4.16. VIRTUALIZATION

Virtualization is now supported on ARM 64

This update introduces support for creating KVM virtual machines on systems that use ARM 64 (also known as AArch64) CPUs. Note, however, that certain virtualization features and functionalities that are available on AMD64 and Intel 64 systems might work differently or be unsupported on ARM 64.

For details, see How virtualization on ARM 64 differs from AMD 64 and Intel 64 .

Jira:RHEL-14097

External snapshots for virtual machines

This update introduces the external snapshot mechanism for virtual machines (VMs), which replaces the previously deprecated internal snapshot mechanism. As a result, you can create, delete, and revert to VM snapshots that are fully supported. External snapshots work more reliably both in the command-line interface and in the RHEL web console. This also applies to snapshots of running VMs, known as live snapshots.

Note, however, that some commands and utilities might still create internal snapshots. To verify that your snapshot is fully supported, ensure that it is configured as **external**. For example:

```
# virsh snapshot-dumpxml VM-name snapshot-name | grep external
<disk name='vda' snapshot='external' type='file'>
```

Jira:RHEL-7528

RHEL now supports Multi-FD migration of virtual machines

With this update, multiple file descriptors (multi-FD) migration of virtual machines is now supported. Multi-FD migration uses multiple parallel connections to migrate a virtual machine, which can speed up the process by utilizing all the available network bandwidth.

It is recommended to use this feature on high-speed networks (20 Gbps and higher).

Jira:RHELDOCS-16970[1]

VM migration now supports post-copy preemption

Post-copy live migrations of virtual machines (VM) now use the **postcopy-preempt** feature, which improves the performance and stability of these migrations.

Jira:RHEL-13004[1], Jira:RHEL-7100

Secure Execution VMs on IBM Z now support cryptographic coprocessors

With this update, you can now assign cryptographic coprocessors as mediated devices to a virtual machine (VM) with IBM Secure Execution on IBM Z.

By assigning a cryptographic coprocessor as a mediated device to a Secure Execution VM, you can now use hardware encryption without compromising the security of the VM.

Jira:RHEL-11597[1]

4th Generation AMD EPYC processors supported on KVM guests

Support for 4th Generation AMD EPYC processors (also known as AMD Genoa) has now been added to the KVM hypervisor and kernel code, and to the libvirt API. This enables KVM virtual machines to use 4th Generation AMD EPYC processors.

Jira:RHEL-7568

New virtualization features in the RHEL web console

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

- Add an SSH public key during virtual machine (VM) creation. This public key will be stored in the **~/.ssh/authorized_keys** file of the designated non-root user on the newly created VM, which provides you with an immediate SSH access to the specified user account.

- Select a **pre-formatted block device** type when creating a new storage pool. This is a more robust alternative to a **physical disk device** type, as it prevents unintentional reformatting of a raw disk device.

This update also changes some default behavior in the Virtual Machines page:

- In the **Add disk** dialog, the **Always attach** option is now set by default.

- The **Create snapshot** action now uses an external snapshot insted of an internal snapshot, which is deprecated in RHEL 9. External snapshots are more reliable and also work for **raw** images, not just for **qcow2** images. You can also select a memory snapshot file location if you want to retain the memory state of the running VM.

Jira:RHELDOCS-17000[1]

virtio-mem is now supported on AMD64 and Intel 64 systems

With this update, RHEL 9 introduces support for the **virtio-mem** feature on AMD64 and Intel 64 systems. With **virtio-mem**, you can dynamically add or remove host memory in virtual machines (VMs).

For more information on **virtio-mem**, see: Adding and removing virtual machine memory by using virtio-mem

Jira:RHELDOCS-17053[1]

You can now replace SPICE with VNC in the web console

With this update, you can use the web console to replace the SPICE remote display protocol with the VNC protocol in an existing virtual machine (VM).

Because the support for the SPICE protocol has been removed in RHEL 9, VMs that use the SPICE protocol fail to start on a RHEL 9 host. For example, RHEL 8 VMs use SPICE by default, so you must switch from SPICE to VNC for a successful migration to RHEL 9.

Jira:RHEL-17434

Improved I/O performance for virtio-blk disk devices

With this update, you can configure a separate IOThread for each virtqueue in a **virtio-blk** disk device. This configuration improves performance for virtual machines with multiple CPUs during intensive I/O workloads.

Jira:RHEL-7416

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

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

Jira:RHEL-7478

4.17. RHEL IN CLOUD ENVIRONMENTS

RHEL instances on EC2 now support IPv6 IMDS connections

With this update, RHEL 8 and 9 instances on Amazon Elastic Cloud Compute (EC2) can use the IPv6 protocol to connect to Instance Metadata Service (IMDS). As a result, you can configure RHEL instances with **cloud-init** on EC2 with a dual-stack IPv4 and IPv6 connection. In addition, you can launch EC2 instances of RHEL with **cloud-init** in IPv6-only subnet.

Jira:RHEL-7278

New cloud-init clean option for deleting generated configuration files

The **cloud-init clean --configs** option has been added for the **cloud-init** utility. You can use this option to delete unnecessary configuration files generated by **cloud-init** on your instance. For example, to delete **cloud-init** configuration files that define network setup, use the following command:

```
cloud-init clean --configs network
```

Jira:RHEL-7311[1]

4.18. CONTAINERS

Podman now supports containers.conf modules

You can use Podman modules to load a predetermined set of configurations. Podman modules are **containers.conf** files in the TOML format.

These modules are located in the following directories, or their subdirectories:

- For rootless users: **$HOME/.config/containers/containers.conf.modules**

- For root users: **/etc/containers/containers.conf.modules**, or **/usr/share/containers/containers.conf.modules**

You can load the modules on-demand with the **podman --module <your_module_name>** command to override the system and user configuration files. Working with modules involve the following facts:

- You can specify modules multiple times by using the **--module** option.

- If **<your_module_name>** is the absolute path, the configuration file will be loaded directly.

- The relative paths are resolved relative to the three module directories mentioned previously.

- Modules in **$HOME** override those in the **/etc/** and **/usr/share/** directories.

For more information, see the upstream documentation.

Jira:RHELPLAN-167829[1]

The Container Tools packages have been updated

The updated Container Tools RPM meta-package, which contain the Podman, Buildah, Skopeo, crun, and runc tools, are now available. Notable bug fixes and enhancements over the previous version include:

Notable changes in Podman v4.9:

- You can now use Podman to load the modules on-demand by using the **podman --module <your_module_name>** command and to override the system and user configuration files.

- A new **podman farm** command with a set of the **create**, **set**, **remove**, and **update** subcommands has been added. With these commands, you can farm out builds to machines running podman for different architectures.

- A new **podman-compose** command has been added, which runs Compose workloads by using an external compose provider such as Docker compose.

- The **podman build** command now supports the **--layer-label** and **--cw** options.

- The **podman generate systemd** command is deprecated. Use Quadlet to run containers and pods under **systemd**.

- The **podman build** command now supports **Containerfiles** with the HereDoc syntax.

- The **podman kube play** command now supports a new **--publish-all** option. Use this option to expose all containerPorts on the host.

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

Jira:RHELPLAN-167796[1]

The Podman v4.9 RESTful API now displays data of progress

With this enhancement, the Podman v4.9 RESTful API now displays data of progress when you pull or push an image to the registry.

Jira:RHELPLAN-167823[1]

Toolbx is now available

With Toolbx, you can install the development and debugging tools, editors, and Software Development Kits (SDKs) into the Toolbx fully mutable container without affecting the base operating system. The Toolbx container is based on the **registry.access.redhat.com/ubi9.4/toolbox:latest** image.

Jira:RHELDOCS-16241[1]

SQLite is now fully supported as a default database backend for Podman

With Podman v4.9, the SQLite database backend for Podman, previously available as Technology Preview, is now fully supported. The SQLite database provides better stability, performance, and consistency when working with container metadata. The SQLite database backend is the default backend for new installations of RHEL 9.4. If you upgrade from a previous RHEL version, the default backend is BoltDB.

If you have explicitly configured the database backend by using the **database_backend** option in the **containers.conf** file, then Podman will continue to use the specified backend.

Jira:RHELPLAN-168180[1]

Administrators can set up isolation for firewall rules by using nftables

You can use Netavark, a Podman container networking stack, on systems without **iptables** installed. Previously, when using the container networking interface (CNI) networking, the predecessor to Netavark, there was no way to set up container networking on systems without **iptables** installed. With this enhancement, the Netavark network stack works on systems with only **nftables** installed and improves isolation of automatically generated firewall rules.

Jira:RHELDOCS-16955[1]

Containerfile now supports multi-line instructions

You can use the multi-line HereDoc instructions (Here Document notation) in the **Containerfile** file to simplify this file and reduce the number of image layers caused by performing multiple **RUN** directives.

For example, the original **Containerfile** can contain the following **RUN** directives:

```
RUN dnf update
RUN dnf -y install golang
RUN dnf -y install java
```

Instead of multiple RUN directives, you can use the HereDoc notation:

```
RUN <<EOF
dnf update
dnf -y install golang
dnf -y install java
EOF
```

Jira:RHELPLAN-168185[1]

The gvisor-tap-vsock package is now available

The **gvisor-tap-vsock** package is an alternative to the **libslirp** user-mode networking library and VPNKit tools and services. It is written in Go and based on the network stack of gVisor. Compared to **libslirp**, the **gvisor-tap-vsock** librarysupports a configurable DNS server and dynamic port forwarding. You can use the **gvisor-tap-vsock** networking library for podman-machine virtual machines. The **podman machine** command for managing virtual machines is currently unsupported on Red Hat Enterprise Linux.

Jira:RHELPLAN-167396[1]

CHAPTER 5. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS

This chapter provides system administrators with a summary of significant changes in the kernel distributed with Red Hat Enterprise Linux 9.4. These changes could include, for example, added or updated **proc** entries, **sysctl**, and **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes.

New kernel parameters

accept_memory=

[MM]

Values:

lazy (default)

By default, unaccepted memory is accepted lazily to avoid prolonged boot times. The lazy option adds some runtime overhead until all memory is eventually accepted. In most cases, the overhead is negligible.

eager

For some workloads or for debugging purposes, you can use **accept_memory=eager** to accept all memory at once during boot.

arm64.nomops

[ARM64]

Unconditionally disable Memory Copy and Memory Set instructions support.

cgroup_favordynmods=

[KNL]

Enable or disable **favordynmods**.

Values:

- **true**

- **false**

Defaults to the value of **CONFIG_CGROUP_FAVOR_DYNMODS**.

early_page_ext

[KNL]

Enforces **page_ext** initialization to earlier stages to cover more early boot allocations.

Note that as side effect, some optimizations might be disabled to achieve that: for example, parallelized memory initialization is disabled. Therefore, the boot process might take longer, especially on systems with much memory.

Available with **CONFIG_PAGE_EXTENSION=y**.

fw_devlink.sync_state=

[KNL]

When all devices that could probe have finished probing, this parameter controls what to do with devices that have not yet received their **sync_state()** calls.

Values:

strict (default)

Continue waiting on consumers to probe successfully.

timeout

Give up waiting on consumers and call **sync_state()** on any devices that have not yet received their **sync_state()** calls after **deferred_probe_timeout** has expired or by **late_initcall()** if **CONFIG_MODULES** is **false**.

ia32_emulation=

[X86-64]

Values:

true

Allows loading 32-bit programs and executing 32-bit syscalls, essentially overriding **IA32_EMULATION_DEFAULT_DISABLED** at boot time.

false

Unconditionally disables IA32 emulation.

kunit.enable=

[KUNIT]

Enable executing KUnit tests. Requires **CONFIG_KUNIT** to be set to be fully enabled.

You can override the default value using **KUNIT_DEFAULT_ENABLED**.

The default is 1 (enabled).

mtrr=debug

[X86]

Enable printing debug information related to MTRR registers at boot time.

rcupdate.rcu_cpu_stall_cputime=

[KNL]

Provide statistics on the CPU time and count of interrupts and tasks during the sampling period. For multiple continuous RCU stalls, all sampling periods begin at half of the first RCU stall timeout.

rcupdate.rcu_exp_stall_task_details=

[KNL]

Print stack dumps of any tasks blocking the current expedited RCU grace period during an expedited RCU CPU stall warning.

spec_rstack_overflow=

[X86]

Control RAS overflow mitigation on AMD Zen CPUs.

Values:

off

 Disable mitigation

microcode

 Enable only microcode mitigation.

safe-ret (default)

 Enable software-only safe RET mitigation.

ibpb

 Enable mitigation by issuing IBPB on kernel entry.

ibpb-vmexit

 Issue IBPB only on VMEXIT. This mitigation is specific to cloud environments.

workqueue.unbound_cpus=

[KNL,SMP]

Specify to constrain one or some CPUs to use in unbound workqueues.

Value: A list of CPUs.

By default, all online CPUs are available for unbound workqueues.

Updated kernel parameters

amd_iommu=

[HW, X86-64]

Pass parameters to the AMD IOMMU driver in the system.

Values:

fullflush

 Deprecated, equivalent to **iommu.strict=1**.

off

 Do not initialize any AMD IOMMU found in the system.

force_isolation

 Force device isolation for all devices. The IOMMU driver is not allowed anymore to lift isolation requirements as needed. This option does not override **iommu=pt**.

force_enable

 Force enable the IOMMU on platforms known to be buggy with IOMMU enabled. Use this option with care.

New: **pgtbl_v1** (default)

 Use version 1 page table for DMA-API.

New: **pgtbl_v2**

Use version 2 page table for DMA-API.

New: irtcachedis

Disable Interrupt Remapping Table (IRT) caching.

nosmt

[KNL, **PPC**, S390]

Disable symmetric multithreading (SMT). Equivalent to **smt=1**.

[KNL, X86, **PPC**]

Disable symmetric multithreading (SMT).

nosmt=force

Force disable SMT. Cannot be undone using the **sysfs** control file.

page_reporting.page_reporting_order=

[KNL]

Minimal page reporting order.

Value: integer.

Adjust the minimal page reporting order.

New: The page reporting is disabled when it exceeds **MAX_ORDER**.

tsc=

Disable clocksource stability checks for TSC.

Values:

[x86] reliable

Mark tsc clocksource as reliable. This disables clocksource verification at runtime, and the stability checks done at bootup. Used to enable high-resolution timer mode on older hardware, and in virtualized environment.

[x86] noirqtime

Do not use TSC to do **irq** accounting. Used to run time disable **IRQ_TIME_ACCOUNTING** on any platforms where RDTSC is slow and this accounting might add overhead.

[x86] unstable

Mark the TSC clocksource as unstable. This marks the TSC unconditionally unstable at bootup and avoids any further wobbles once the TSC watchdog notices.

[x86] nowatchdog

Disable clocksource watchdog. Used in situations with strict latency requirements, where interruptions from clocksource watchdog are not acceptable.

[x86] recalibrate

Force recalibration against a HW timer (HPET or PM timer) on systems whose TSC frequency was obtained from HW or FW using either an MSR or CPUID(0x15). Warn if the difference is more than 500 ppm.

New: [x86] watchdog

Use TSC as the watchdog clocksource with which to check other HW timers (HPET or PM timer), but only on systems where TSC has been deemed trustworthy.
An earlier **tsc=nowatchdog** suppresses this. A later **tsc=nowatchdog** overrides this. A console message flags any such suppression or overriding.

usbcore.authorized_default=

[USB]

Default USB device authorization.

Values:

New: -1 (default)

Authorized (same as 1).

0

Not authorized.

1

Authorized.

2

Authorized if the device connects to an internal port.

Removed kernel parameters

- **cpu0_hotplug**

- **sysfs.deprecated**

New sysctl parameters

io_uring_group

Values:

1

A process must either be privileged (**CAP_SYS_ADMIN**) or be in the **io_uring_group** group to create an **io_uring** instance.

-1 (default)

Only processes with the **CAP_SYS_ADMIN** capability can create **io_uring** instances.

numa_balancing_promote_rate_limit_MBps

Too high promotion or demotion throughput between different memory types might hurt application latency. You can use this parameter to rate-limit the promotion throughput. The per-node maximum promotion throughput in MB/s is limited to be no more than the set value.

A rule of thumb is to set this to less than 1/10 of the PMEM node write bandwidth.

Updated sysctl parameters

io_uring_disabled

Prevents all processes from creating new **io_uring** instances. Enabling this shrinks the attack surface of the kernel.

Values:

New: 0

All processes can create **io_uring** instances as normal.

New: 1

io_uring creation is disabled for unprivileged processes not in the io_uring_group group.
io_uring_setup() fails with **-EPERM**. Existing **io_uring** instances can still be used.
See the documentation for **io_uring_group** for more information.

New: 2 (default)

io_uring creation is disabled for all processes. **io_uring_setup()** always fails with **-EPERM**. Existing
io_uring instances can still be used.

CHAPTER 6. DEVICE DRIVERS

6.1. NEW DRIVERS

Table 6.1. Cryptographic drivers

Description	Name	Limited to architectures
IAA Compression Accelerator Crypto Driver	iaa_crypto	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	intel_qat	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_4xxx	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_c3xxx	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_c3xxxvf	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_c62x	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_c62xvf	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_dh895xcc	AMD and Intel 64-bit architectures
Intel® QuickAssist Technology - 0.6.0	qat_dh895xccvf	AMD and Intel 64-bit architectures

Table 6.2. Network drivers

Description	Name	Limited to architectures
	bcm-phy-ptp	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	mt7925-common	64-bit ARM architecture, AMD and Intel 64-bit architectures
	mt7925e	64-bit ARM architecture, AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
	mt792x-lib	64-bit ARM architecture, AMD and Intel 64-bit architectures
CAN bus driver for Bosch M_CAN controller on PCI bus	m_can_pci	IBM Power Systems, AMD and Intel 64-bit architectures
CAN bus driver for Bosch M_CAN controller	m_can	IBM Power Systems, AMD and Intel 64-bit architectures
CAN driver for 8 devices USB2CAN interfaces	usb_8dev	IBM Power Systems, AMD and Intel 64-bit architectures
CAN driver for EMS Dr. Thomas Wuensche CAN/USB interfaces	ems_usb	IBM Power Systems, AMD and Intel 64-bit architectures
CAN driver for Kvaser CAN/USB devices	kvaser_usb	IBM Power Systems, AMD and Intel 64-bit architectures
CAN driver for PEAK-System USB adapters	peak_usb	IBM Power Systems, AMD and Intel 64-bit architectures
Intel® Infrastructure Data Path Function Linux Driver	idpf	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Marvell 88Q2XXX 100/1000BASE-T1 Automotive Ethernet PHY driver	marvell-88q2xxx	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Marvell Octeon EndPoint NIC Driver	octeon_ep	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Microchip 251x/25625 CAN driver	mcp251x	AMD and Intel 64-bit architectures
Microchip MCP251xFD Family CAN controller driver	mcp251xfd	AMD and Intel 64-bit architectures
NXP imx8 DWMAC Specific Glue layer	dwmac-imx	64-bit ARM architecture
	bcm-phy-ptp	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
Realtek 802.11ax wireless 8852C driver	rtw89_8852c	64-bit ARM architecture, AMD and Intel 64-bit architectures
Realtek 802.11ax wireless 8852CE driver	rtw89_8852ce	64-bit ARM architecture, AMD and Intel 64-bit architectures
serial line CAN interface	slcan	IBM Power Systems, AMD and Intel 64-bit architectures
Socket-CAN driver for PEAK PCAN PCIe/M.2 FD family cards	peak_pciefd	IBM Power Systems, AMD and Intel 64-bit architectures
	bcm-phy-ptp	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	mt7925-common	64-bit ARM architecture, AMD and Intel 64-bit architectures
	mt7925e	64-bit ARM architecture, AMD and Intel 64-bit architectures
	mt792x-lib	64-bit ARM architecture, AMD and Intel 64-bit architectures

Table 6.3. Platform drivers

Description	Name	Limited to architectures
AMD HSMP Platform Interface Driver - 2.0	amd_hsmp	AMD and Intel 64-bit architectures
AMD Platform Management Framework Driver	amd-pmf	AMD and Intel 64-bit architectures
Intel TPMI enumeration module	intel_vsec_tpmi	AMD and Intel 64-bit architectures
Intel TPMI SST Driver	isst_tpmi	AMD and Intel 64-bit architectures
Intel TPMI UFS Driver	intel-uncore-frequency-tpmi	AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
Intel Uncore Frequency Common Module	intel-uncore-frequency-common	AMD and Intel 64-bit architectures
Intel Uncore Frequency Limits Driver	intel-uncore-frequency	AMD and Intel 64-bit architectures
Intel WMI Thunderbolt force power driver	intel-wmi-thunderbolt	AMD and Intel 64-bit architectures
Mellanox PMC driver	mlxbf-pmc	64-bit ARM architecture
	intel-hid	AMD and Intel 64-bit architectures
	isst_tpmi_core	AMD and Intel 64-bit architectures

Table 6.4. Graphics drivers and miscellaneous drivers

Description	Name	Limited to architectures
AMD XCP Platform Devices	amdxcp	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
DRM execution context	drm_exec	
Range suballocator helper	drm_suballoc_helper	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	regmap-ram	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	regmap-raw-ram	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	regmap-ram	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
	regmap-raw-ram	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	regmap-ram	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
	regmap-raw-ram	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Arm FF-A interface driver	ffa-module	64-bit ARM architecture
NVIDIA BlueField-3 GPIO Driver	gpio-mlxbf3	64-bit ARM architecture
I/O Address Space Management for passthrough devices	iommufd	
CS42L43 Core Driver	cs42l43	AMD and Intel 64-bit architectures
CS42L43 SoundWire Driver	cs42l43-sdw	AMD and Intel 64-bit architectures
MEI GSC Proxy	mei_gsc_proxy	AMD and Intel 64-bit architectures
	pwrseq_emmc	64-bit ARM architecture
	pwrseq_simple	64-bit ARM architecture
SDHCI platform driver for Synopsys DWC MSHC	sdhci-of-dwcmshc	64-bit ARM architecture
	arm_cspmu_module	64-bit ARM architecture
NVIDIA pinctrl driver	pinctrl-mlxbf3	64-bit ARM architecture
NXP i.MX93 power domain driver	imx93-pd	64-bit ARM architecture
Intel RAPL TPMI Driver	intel_rapl_tpmi	AMD and Intel 64-bit architectures

Description	Name	Limited to architectures
Mellanox BlueField power driver	pwr-mlxbf	64-bit ARM architecture
NXP i.MX93 src driver	imx93-src	64-bit ARM architecture
Provide Trusted Security Module attestation reports via configfs	tsm	AMD and Intel 64-bit architectures

6.2. UPDATED DRIVERS

Table 6.5. Storage driver updates

Description	Name	Current version	Limited to architectures
Broadcom MegaRAID SAS Driver	megaraid_sas	07.727.03.00-rc1	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Driver for Microchip Smart Family Controller	smartpqi	2.1.24-046	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
Emulex LightPulse Fibre Channel SCSI driver	lpfc	0:14.2.0.16	64-bit ARM architecture, IBM Power Systems, AMD and Intel 64-bit architectures
MPI3 Storage Controller Device Driver	mpi3mr	8.5.0.0.50	

CHAPTER 7. AVAILABLE BPF FEATURES

This chapter provides the complete list of Berkeley Packet Filter (BPF) features available in the kernel of this minor version of Red Hat Enterprise Linux 9. The tables include the lists of:

- System configuration and other options

- Available program types and supported helpers

- Available map types

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

Table 7.1. System configuration and other options

Option	Value
unprivileged_bpf_disabled	2 (bpf() syscall restricted to privileged users, admin can change)
JIT compiler	1 (enabled)
JIT compiler hardening	1 (enabled for unprivileged users)
JIT compiler kallsyms exports	1 (enabled for root)
Memory limit for JIT for unprivileged users	528482304
CONFIG_BPF	y
CONFIG_BPF_SYSCALL	y
CONFIG_HAVE_EBPF_JIT	y
CONFIG_BPF_JIT	y
CONFIG_BPF_JIT_ALWAYS_ON	y
CONFIG_DEBUG_INFO_BTF	y
CONFIG_DEBUG_INFO_BTF_MODULES	y
CONFIG_CGROUPS	y
CONFIG_CGROUP_BPF	y
CONFIG_CGROUP_NET_CLASSID	y
CONFIG_SOCK_CGROUP_DATA	y

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

Option	Value
CONFIG_TEST_BPF	m
CONFIG_HZ	1000
bpf() syscall	available
Large program size limit	available
Bounded loop support	available
ISA extension v2	available
ISA extension v3	available

Table 7.2. Available program types and supported helpers

Program type	Available helpers
socket_filter	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_load_bytes_relative, bpf_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_wri table	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_strtol, bpf_strtoul, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_copy_from_user, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_copy_from_user_task, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
cgroup_sockopt	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_tcp_sock, bpf_strtol, bpf_strtoul, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_get_retval, bpf_set_retval, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
tracing	
struct_ops	
ext	
lsm	

Program type	Available helpers
sk_lookup	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_strtol, bpf_strtoul, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_skc_to_unix_sock, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete
syscall	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_get_socket_cookie, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_strtol, bpf_strtoul, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_send_signal, bpf_skb_output, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_xdp_output, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_get_task_stack, bpf_d_path, bpf_copy_from_user, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_sock_from_file, bpf_for_each_map_elem, bpf_snprintf, bpf_sys_bpf, bpf_btf_find_by_name_kind, bpf_sys_close, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_get_func_ip, bpf_task_pt_regs, bpf_get_branch_snapshot, bpf_skc_to_unix_sock, bpf_kallsyms_lookup_name, bpf_find_vma, bpf_loop, bpf_strncmp, bpf_xdp_get_buff_len, bpf_copy_from_user_task, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_skc_to_mptcp_sock, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

Program type	Available helpers
netfilter	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_strtol, bpf_strtoul, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_for_each_map_elem, bpf_snprintf, bpf_timer_init, bpf_timer_set_callback, bpf_timer_start, bpf_timer_cancel, bpf_task_pt_regs, bpf_loop, bpf_strncmp, bpf_kptr_xchg, bpf_map_lookup_percpu_elem, bpf_dynptr_from_mem, bpf_ringbuf_reserve_dynptr, bpf_ringbuf_submit_dynptr, bpf_ringbuf_discard_dynptr, bpf_dynptr_read, bpf_dynptr_write, bpf_dynptr_data, bpf_ktime_get_tai_ns, bpf_user_ringbuf_drain, bpf_cgrp_storage_get, bpf_cgrp_storage_delete

Table 7.3. Available map types

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

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

CHAPTER 8. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 9.4 that have a significant impact on users.

8.1. INSTALLER AND IMAGE CREATION

Anaconda displays WWID identifiers for multipath storage devices on the Installation Destination screen

Previously, Anaconda did not display any details, for example, device number, WWPN, or LUN for the multipath storage devices. As a consequence, it was difficult to select the correct installation destination from the **Installation Destination** > **Add a disk** screen. With this update, Anaconda now displays WWID identifiers for multipath storage devices. As a result, you can now easily identify and select the desired installation destination on the advanced storage device screen.

Jira:RHEL-11384[1]

Installer now accepts additional time zone definitions in Kickstart files

Anaconda switched to a different, more restrictive method of validating time zone selections. This caused some time zone definitions, such as Japan, to be no longer valid despite being accepted in previous versions. Legacy Kickstart files with these definitions had to be updated. Otherwise, they would default to the **Americas/New_York time** zone.

The list of valid time zones was previously taken from **pytz.common_timezones** in the **pytz** Python library. This update changes the validation settings for the **timezone** Kickstart command to use **pytz.all_timezones**, which is a superset of the **common_timezones** list, and allows significantly more time zones to be specified. This change ensures that old Kickstart files made for Red Hat Enterprise Linux 6 still specify valid time zones.

Note: This change only applies to the **timezone** Kickstart command. The time zone selection in the graphical and text-based interactive interfaces remains unchanged. Existing Kickstart files for Red Hat Enterprise Linux 9 that had valid time zone selections do not require any updates.

Jira:RHEL-13150[1]

The installer now correctly creates bond device with multiple ports and a BOOTIF option

Previously, the installer created incorrect connection profiles when the installation was booted with a bond network device with multiple ports along with the **BOOTIF** boot option. Consequently, the device used by the BOOTIF option was not added to the bond device though it was configured as one of its ports.

With this update, the installer now correctly creates profiles in **initramfs** when the **BOOTIF** boot option is used. As a result, all the specified ports are now added to the bond device on the installed system.

Jira:RHEL-4766

Anaconda replaces the misleading error message when failing to boot an installation image

Previously, when the installation program failed to boot the installation image, for example due to missing source of **stage2** specified in **inst.stage2** or **inst.repo**, Anaconda displayed the following misleading error message:

> /run/anaconda/initrd_errors.txt: No such file or directory

With this update, Anaconda issues a proper warning message to minimize the confusion.

Jira:RHEL-5638

The new version of **xfsprogs** no longer shrinks the size of **/boot**

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

Jira:RHEL-7999

8.2. SECURITY

Libreswan accepts IPv6 SAN extensions

Previously, IPsec connection failed when setting up certificate-based authentication with a certificate that contained a subjectAltName (SAN) extension with an IPv6 address. With this update, the **pluto** daemon has been modified to accept IPv6 SAN as well as IPv4. As a result, IPsec connection is now correctly established with IPv6 address embedded in the certificate as an ID.

Jira:RHEL-12278

Rules for managing virtual routing with **ip vrf** are added to the SELinux policy

You can use the **ip vrf** command to manage virtual routing of other network services. Previously, **selinux-policy** did not contain rules to support this usage. With this update, SELinux policy rules allow explicit transitions from the **ip** domain to the **httpd**, **sshd**, and **named** domains. These transitions apply when the **ip** command uses the **setexeccon** library call.

Jira:RHEL-14246[1]

SELinux policy denies SSH login for unconfined users when **unconfined_login** is set to **off**

Previously, the SELinux policy was missing a rule to deny unconfined users to log in via SSH when the **unconfined_login** boolean was set to **off**. As a consequence, with **unconfined_login** set to **off**, users still could log in with SSHD to an unconfined domain. This update adds a rule to the SELinux policy, and as a result, users cannot log in via **sshd** as unconfined when **unconfined_login** is **off**.

Jira:RHEL-1551

SELinux policy allows **rsyslogd** to execute confined commands

Previously, the SELinux policy was missing a rule to allow the **rsyslogd** daemon to execute SELinux-confined commands, such as **systemctl**. As a consequence, commands executed as an argument of the **omprog** directive failed. This update adds rules to the SELinux policy so that executables in the **/usr/libexec/rsyslog** directory that are run as an argument of **omprog** are in the **syslogd_unconfined_script_t** unconfined domain. As a result, commands executed as an argument of **omprog** finish successfully.

Jira:RHEL-11174

SELinux policy allows **rsyslogd** to execute confined commands

Previously, the SELinux policy was missing a rule to allow the **rsyslogd** daemon to execute SELinux-confined commands, such as **systemctl**. As a consequence, commands executed as an argument of the **omprog** directive failed. This update adds rules to the SELinux policy so that executables in the

/usr/libexec/rsyslog directory that are run as an argument of **omprog** are in the **syslogd_unconfined_script_t** unconfined domain. As a result, commands executed as an argument of **omprog** finish successfully.

Jira:RHEL-10087

kmod runs in the SELinux MLS policy

Previously, the SELinux did not assign a private type for the **/var/run/tmpfiles.d/static-nodes.conf** file. As a consequence, the **kmod** utility may fail to work in the SELinux multi-level security (MLS) policy. This update adds the **kmod_var_run_t** label for **/var/run/tmpfiles.d/static-nodes.conf** to the SELinux policy, and as a result, **kmod** runs successfully in the SELinux MLS policy.

Jira:RHEL-1553

selinux-autorelabel runs in SELinux MLS policy

Previously, the SELinux policy did not assign a private type for the **/usr/libexec/selinux/selinux-autorelabel** utility. As a consequence, **selinux-autorelabel.service** might fail to work in the SELinux multi-level security (MLS) policy. This update adds the **semanage_exec_t** label to **/usr/libexec/selinux/selinux-autorelabel**, and as a result, **selinux-autorelabel.service** runs successfully in the SELinux MLS policy.

Jira:RHEL-14289

/bin = /usr/bin file context equivalency rule added to SELinux policy

Previously, the SELinux policy did not contain the **/bin = /usr/bin** file context equivalency rule. As a consequence, the **restorecond** daemon did not work correctly. This update adds the missing rule to the policy, and as a consequence, **restorecond** works correctly in SELinux enforcing mode.

> **IMPORTANT**
>
> This change overrides any local policy modules which use file context specification for a pattern in **/bin**.

Jira:RHEL-5032

SELinux policy contains rules for additional services and applications

This version of the **selinux-policy** package contains additional rules. Most notably, users in the **sysadm_r** role can execute the following commands:

- **sudo traceroute** (RHEL-14077)

- **sudo tcpdump** (RHEL-15432)

Jira:RHEL-15432

SELinux policy adds permissions for QAT firmware

Previously, when updating the Intel QuickAssist Technology (QAT) with the **Intel VT-d** kernel option enabled, missing SELinux permissions caused denials. This update adds additional permissions for the **qat** service. As a result, QAT can be updated correctly.

Jira:RHEL-19051[1]

Rsyslog can execute privileged commands through **omprog**

Previously, the **omprog** module of Rsyslog could not execute certain external programs, especially programs that contain privileged commands. As a consequence, the use of scripts that involve privileged commands through **omprog** was restricted. With this update, the SELinux policy was adjusted. Place your scripts into the **/usr/libexec/rsyslog** directory to ensure compatibility with the adjusted SELinux policy. As a result, Rsyslog now can execute scripts, including those with privileged commands, through the **omprog** module.

Jira:RHEL-5196

The **semanage fcontext** command no longer reorders local modifications

The **semanage fcontext -l -C** command lists local file context modifications stored in the **file_contexts.local** file. The **restorecon** utility processes the entries in the **file_contexts.local** from the most recent entry to the oldest. Previously, **semanage fcontext -l -C** listed the entries in an incorrect order. This mismatch between processing order and listing order caused problems when managing SELinux rules. With this update, **semanage fcontext -l -C** displays the rules in the correct and expected order, from the oldest to the newest.

Jira:RHEL-25263[1]

CardOS 5.3 cards with offsets no longer cause problems in OpenSC

Previously, file caching did not work correctly for some CardOS 5.3 cards that stored certificates on different offsets of a single PKCS #15 file. This occurred because file caching ignored the offset part of the file, which caused repetitive overriding of the cache and reading invalid data from file cache. The problem was identified and fixed upstream, and after this update, CardOS 5.3 cards work correctly with the file cache.

Jira:RHEL-4079[1]

8.3. SUBSCRIPTION MANAGEMENT

subscription-manager no longer retains nonessential text in the terminal

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

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

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

Bugzilla:2136694[1]

8.4. SOFTWARE MANAGEMENT

The **librhsm** library now returns the correct **/etc/rhsm-host** prefix if **librhsm** is run in a container

The **librhsm** library rewrites path prefixes to CA certificates from the **/etc/rhsm** to **/etc/rhsm-host** path if **librhsm** is run in a container. Previously, **librhsm** returned the wrong **/etc/rhsm-host-host** prefix because of a string manipulation mistake. With this update, the issue has been fixed, and the **librhsm** library now returns the correct **/etc/rhsm-host** prefix.

Jira:RHEL-14224

systemd now correctly manages the /run/user/0 directory created by librepo

Previously, if the **librepo** functions were called from an Insights client before logging in root, the **/run/user/0** directory could be created with a wrong SELinux context type. This prevented **systemd** from cleaning the directory after you logged out from root.

With this update, the **librepo** package now sets a default creation type according to default file system labeling rules defined in a SELinux policy. As a result, **systemd** now correctly manages the **/run/user/0** directory created by **librepo**.

Jira:RHEL-11240

systemd now correctly manages the /run/user/0 directory created by libdnf

Previously, if the **libdnf** functions were called from an Insights client before logging in root, the **/run/user/0** directory could be created with a wrong SELinux context type. This prevented **systemd** from cleaning the directory after you logged out from root.

With this update, the **libdnf** package now sets a default creation type according to default file system labeling rules defined in a SELinux policy. As a result, **systemd** now correctly manages the **/run/user/0** directory created by **libdnf**.

Jira:RHEL-11238

The dnf needs-restarting --reboothint command now recommends a reboot to update the CPU microcode

To fully update the CPU microcode, you must reboot a system. Previously, when you installed the **microcode_ctl** package, which contains the updated CPU microcode, the **dnf needs-restarting --reboothint** command did not recommend the reboot. With this update, the issue has been fixed, and **dnf needs-restarting --reboothint** now recommends a reboot to update the CPU microcode.

Jira:RHEL-4600

8.5. SHELLS AND COMMAND-LINE TOOLS

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

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

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

Jira:RHEL-16278

ReaR now determines the presence of a BIOS bootloader when both BIOS and UEFI bootloaders are installed

Previously, in a hybrid bootloader setup (**UEFI** and **BIOS**), when **UEFI** was used to boot, Relax-and-Recover (ReaR) restored only the **UEFI** bootloader and not the **BIOS** bootloader. This would result in a system that had a **GUID Partition Table** (**GPT**), a BIOS Boot Partition, but not a BIOS bootloader. In this situation, ReaR failed to create the rescue image, the attempt to produce a backup or a rescue image by using the **rear mkbackup** or **rear mkrescue** command would fail with the following error message:

> ERROR: Cannot autodetect what is used as bootloader, see default.conf about 'BOOTLOADER'.

With this update, ReaR determines the presence of both **UEFI** and **BIOS** bootloaders, restores them, and does not fail when it does not encounter the **BIOS** bootloader on the system with the **BIOS** Boot Partition in **GPT**. As a result, systems with the hybrid **UEFI** and **BIOS** bootloader setup can be backed up and recovered multiple times.

Jira:RHEL-16864[1]

ReaR no longer uses the logbsize, sunit and swidth mount options during recovery

Previously, when restoring an **XFS** file system with the parameters different from the original ones by using the **MKFS_XFS_OPTIONS** configuration setting, Relax-and-Recover (ReaR) mounted this file system with mount options applicable for the original file system, but not for the restored file system. As a consequence, the disk layout recreation would fail with the following error message when ReaR ran the **mount** command :

> wrong fs type, bad option, bad superblock on and missing codepage or helper program, or other error.

The kernel log displayed either of the following messages:

> logbuf size must be greater than or equal to log stripe size

> alignment check failed: sunit/swidth vs. agsize

With this update, ReaR avoids using the **logbsize**, **sunit** and **swidth** mount options when mounting recreated **XFS** file systems. As a result, when you use the **MKFS_XFS_OPTIONS** configuration setting, the disk layout recreation succeeds.

Jira:RHEL-10478[1]

ReaR recovery no longer fails on systems with a small thin pool metadata size

Previously, ReaR did not save the size of the pool metadata volume when saving a layout of an LVM volume group with a thin pool. During recovery, ReaR recreated the pool with the default size even if the system used a non-default pool metadata size.

As a consequence, when the original pool metadata size was smaller than the default size and no free space was available in the volume group, the layout recreation during system recovery failed with a message in the log similar to these examples:

> Insufficient free space: 230210 extents needed, but only 230026 available

or

> Volume group "vg" has insufficient free space (16219 extents): 16226 required.

With this update, the recovered system has a metadata volume with the same size as the original system. As a result, the recovery of a system with a small thin pool metadata size and no extra free space in the volume group finishes successfully.

Jira:RHEL-6984

ReaR now preserves logs from the **bprestore** command of NetBackup in the rescue system and the recovered system

Previously, when using the NetBackup integration (**BACKUP=NBU**), ReaR added the log from the **bprestore** command during recovery to a directory that was deleted on exit. Additionally, ReaR did not save further logs produced by the command under the **/usr/openv/netbackup/logs/bprestore/** directory on the recovered system.

As a consequence, if the **bprestore** command failed during recovery, the logs were deleted unless the **rear recover** command was run with the **-d** or **-D** option. Moreover, even if the recovery finished successfully, the logs under **/usr/openv/netbackup/logs/bprestore/** directory were lost after a reboot and could not be examined.

With this update, ReaR keeps the log from the **bprestore** command in the **/var/lib/rear/restore** directory in the rescue system where it persists after the **rear recover** command has finished until the rescue system is rebooted. If the system is recovered, all logs from **/usr/openv/netbackup/logs/bprestore/** are copied to the **/var/log/rear/recover/restore** directory together with the log from **/var/lib/rear/restore** in case further examination is required.

Jira:RHEL-17393

ReaR no longer fails during recovery if the **TMPDIR** variable is set in the configuration file

Previously, the ReaR default configuration file **/usr/share/rear/conf/default.conf** contained the following instructions:

```
# To have a specific working area directory prefix for Relax-and-Recover
# specify in /etc/rear/local.conf something like
#
#   export TMPDIR="/prefix/for/rear/working/directory"
#
# where /prefix/for/rear/working/directory must already exist.
# This is useful for example when there is not sufficient free space
# in /tmp or $TMPDIR for the ISO image or even the backup archive.
```

The instructions mentioned above did not work correctly because the **TMPDIR** variable had the same value in the rescue environment, which was not correct if the directory specified in the **TMPDIR** variable did not exist in the rescue image.

As a consequence, when the rescue image was booted, setting and exporting **TMPDIR** in the **/etc/rear/local.conf** file led to the following error :

```
mktemp: failed to create file via template '/prefix/for/rear/working/directory/tmp.XXXXXXXXXX': No
such file or directory
cp: missing destination file operand after '/etc/rear/mappings/mac'
Try 'cp --help' for more information.
No network interface mapping is specified in /etc/rear/mappings/mac
```

or the following error and abort later, when running **rear recover**:

> ERROR: Could not create build area

With this update, ReaR unsets the **TMPDIR** variable in the rescue environment. ReaR also detects when the variable has been set in **/etc/rear/local.conf**, and prints a warning if the variable is set. The comment in **/usr/share/rear/conf/default.conf** has been changed to instruct to set and export **TMPDIR** in the environment before executing **rear** instead of setting it in **/etc/rear/local.conf**.

If the command **export TMPDIR=…** is used in **/etc/rear/local.conf**, ReaR now prints the following warning:

> Warning: Setting TMPDIR in a configuration file is deprecated. To specify a working area directory prefix, export TMPDIR before executing 'rear'

As a result, the recovery is successful in the described configuration.

Setting **TMPDIR** in a configuration file like **/etc/rear/local.conf** is now deprecated and the functionality will be removed in a future release. It is recommended to remove such settings from **/etc/rear/local.conf**, and to set and export **TMPDIR** in the environment before calling ReaR instead.

Jira:RHEL-24847

8.6. NETWORKING

wwan_hwsim is now in the kernel-modules-internal package

The **wwan_hwsim** kernel module provides a framework for simulating and testing various networking scenarios that use wireless wide area network (WWAN) devices. Previously, **wwan_hwsim** was a part of the **kernel-modules-extra** package. However, with this release, it is moved to the **kernel-modules-internal** package, which contains other similarly-oriented utilities. Note that the WWAN feature for PCI modem is still a Technology Preview.

Jira:RHEL-24618[1]

The xdp-loader features command now works as expected

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

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

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

Jira:RHEL-3382

Mellanox ConnectX-5 adapter works in the DMFS mode

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

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

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

Jira:RHEL-9897[1]

8.7. KERNEL

crash was rebased to version 8.0.4

The **crash** utility was upgraded to version 8.0.4, which provides multiple bugfixes. Notable repairs include:

- Fixed the segmentation fault when the non-panicking CPUs failed to stop during the kernel panic.

- The critical error incorrectly did not cause the kernel panic when the **panic_on_oops** kernel parameter was disabled.

- The **crash** utility did not properly resolve the hashed freelist pointers for the kernels compiled with the **CONFIG_SLAB_FREELIST_HARDENED=y** configuration option.

- A change in the kernel module memory layout terminology. The change replaced **module_layout** with **module_memory** to better indicate memory-related aspects of the **crash** utility. Without this change, **crash** cannot start a session with an error message like this:

  > crash: invalid structure member offset: module_core_size
  >   FILE: kernel.c LINE: 3787 FUNCTION: module_init()

Jira:RHEL-9009

tuna launches GUI when needed

Previously, if you ran the **tuna** utility without any subcommand, it would launch the GUI. This behavior was desirable if you had a display. In the opposite case, **tuna** on a machine without a display would not exit gracefully. With this update, **tuna** detects whether you have a display, and the GUI is launched or not launched accordingly.

Jira:RHEL-8859[1]

Intel TPM chips are now detected correctly

Previously, a side effect in a bug fix to AMD Trusted Platform Module (TPM) chips also affected Intel TPM chips. As a consequence, RHEL failed to detect certain Intel TPM chips.

With this update, the AMD TPM bug fix has been revised. As a result, RHEL now detects the Intel TPM chips correctly.

Jira:RHEL-18985[1]

RHEL previously failed to recognize NVMe disks when VMD was enabled

When you reset or reattached a driver, the Volume Management Device (VMD) domain previously did not soft-reset. Consequently, the hardware could not properly detect and enumerate its devices. With this update, the operating system with VMD enabled now correctly recognizes NVMe disks, especially when resetting a server or working with a VM machine.

Bugzilla:2128610[1]

8.8. FILE SYSTEMS AND STORAGE

multipathd now successfully removes devices that have outstanding queued I/O

Previously, the **multipathd** command did not disable the **queue_if_no_path** parameter before removing a device. This was possible only if there was an outstanding queued I/O to the multipath device itself, and not to the partition devices. Consequently, **multipathd** would hang, and could no longer maintain the multipath devices. With this update, the **multipathd** now disables queuing before executing the remove command such as **multipath -F**, **multipath -f <device>**, **multipathd remove maps**, or **multipathd remove map <device>**. As a result, **multipathd** now successfully removes devices that have outstanding queued I/O.

Jira:RHEL-4998[1]

The no_read_workqueue, no_write_workqueue, and try_verify_in_taskle options of the dm-crypt and dm-verity devices are temporarily disabled

Previously, the **dm-crypt** devices created by using either the **no_read_workqueue** or **no_write_workqueue** option and **dm-verity** devices created by using the **try_verify_in_tasklet** option caused memory corruption. Consequently, random kernel memory was corrupted, which caused various system problems. With this update, these options are temporarily disabled. Note that this fix can cause **dm-verity** and **dm-crypt** to perform slower on some workloads.

Jira:RHEL-23572[1]

Multipathd now checks if a device is incorrectly queuing I/O

Previously, a multipath device restarted queuing I/O, even though it was configured to fail, under the following conditions:

- The multipath device was configured with the **queue_if_no_paths** parameter set to a number of retries.

- A path device was removed from the multipath device that had no working paths and was no longer queuing I/O.

With this update, the issue has been fixed. As a result, multipath devices no longer restarts queuing I/O if the queuing is disabled and a path is removed while there are no usable paths.

Jira:RHEL-17234[1]

Removing duplicate entry from nvmf_log_connect_error

Previously, due to a duplicate commit merge error, a log message was repeated in the **nvmf_log_connect_error** kernel function. Consequently, when the kernel was unable to connect to a fabric-attached Non-volatile Memory Express (NVMe) device, the **Connect command failed** message appeared twice. With this update, the duplicate log message is now removed from the kernel, resulting in only a single log message available for each error.

Jira:RHEL-21545[1]

The kernel no longer crashes when namespaces are added and removed

Previously, when NVMe namespaces were rapidly added and removed, a namespace disappeared between successive commands used to probe the namespace. In a specific case, a storage array did not

return an **invalid namespace** error but instead returned a buffer filled with zero. Consequently, the kernel crashed due to the **divide-by-zero** error. With this update, the kernel now validates data from responses to both the Identify Namespace data structure issued to the storage. As a result, the kernel no longer crashes.

Jira:RHEL-14751[1]

The newly allocated sections of the data device are now properly aligned

Previously, when a Stratis pool was expanded, it was possible to allocate the new regions of the pool. But the newly allocated regions were not correctly aligned with the previously allocated regions. Consequently, it could cause a performance degradation along with a non-zero entry in the Stratis thin pool's **alignment_offset** file in **sysfs**. With this update, when the pool expands, the newly allocated region of the data device is properly aligned with the previously allocated region. As a result, there is no degradation in performance and no non-zero entry in the Stratis thin pool's **alignment_offset** file in **sysfs**.

Jira:RHEL-16736

System boots correctly when adding a NVMe-FC device as a mount point in /etc/fstab

Previously, due to a known issue in the **nvme-cli nvmf-autoconnect systemd** services, systems failed to boot while adding the Non-volatile Memory Express over Fibre Channel (NVMe-FC) devices as a mount point in the **/etc/fstab** file. Consequently, the system entered into an emergency mode. With this update, a system boots without any issue when mounting an NVMe-FC device.

Jira:RHEL-8171[1]

LUNs are now visible during the operating system installation

Previously, the system was not using the authentication information from firmware sources, specifically in cases involving iSCSI hardware offload with CHAP (Challenge-Handshake Authentication Protocol) authentication stored in the iSCSI iBFT (Boot Firmware Table). As a consequence, the iSCSI login failed during installation.

With the fix in the **udisks2-2.9.4-9.el9** firmware authentication, this issue is now resolved and LUNs are visible during the installation and initial boot.

Bugzilla:2213769[1]

8.9. HIGH AVAILABILITY AND CLUSTERS

Configuring the **tls** and **keep_active_partition_tie_breaker** quorum device options without specifying **--force**

Previously, when configuring a quorum device, a user could not configure the **tls** and **keep_active_partition_tie_breaker** options for a quorum device model **net** without specifying the **--force** option. With this update, configuring these options no longer requires you to specify   **--force**.

Jira:RHEL-7746

Issues with moving and banning clone and bundle resources now corrected

This bug fix addresses two limitations of moving bundled and clone resources:

- When a user tried to move a bundled resource out of its bundle or ban it from running in its bundle, **pcs** created a constraint but the constraint had no effect. This caused the move to fail

with an error message. With this fix, **pcs** disallows moving and banning bundled resources from their bundles and prints an error message noting that bundled resources cannot be moved out of their bundles.

- When a user tried to move a bundle or clone resource, **pcs** exited with an error message noting that bundle or clone resources cannot be moved. This fix relaxes validation of move commands. It is now possible to move clone and bundle resources. When moving clone resources, you must specify a destination node if more than one instance of a clone is running. Only one-replica bundles can be moved.

Jira:RHEL-7744

Output of `pcs status` command no longer shows warning for expired constraints

Previously, when moving a cluster resource created a temporary location constraint, the **pcs status** command displayed a warning even after the constraint expired. With this fix, the **pcs status** command filters out expired constraints and they no longer generate a warning message in the command output.

Jira:RHEL-7669

Disabling the `auto_tie_breaker` quorum option no longer allowed when SBD fencing requires it

Previously, **pcs** allowed a user to disable the **auto_tie_breaker** quorum option even when a cluster configuration required this option for SBD fencing to work correctly. With this fix, **pcs** generates an error message when a user attempts to disable **auto_tie_breaker** on a system where SBD fencing requires that the **auto_tie_breaker** option be enabled.

Jira:RHEL-7730

8.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

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

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

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

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

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

Jira:RHEL-6600

8.11. COMPILERS AND DEVELOPMENT TOOLS

ldconfig no longer crashes after an interrupted system upgrade

Previously, the **ldconfig** utility terminated unexpectedly with a segmentation fault when processing incomplete shared objects left in the **/usr/lib64** directory after an interrupted system upgrade. With this update, **ldconfig** ignores temporary files written during system upgrades. As a result, **ldconfig** no longer crashes after an interrupted system upgrade.

Jira:RHEL-14383

glibc now uses the number of configured processors for **malloc** arena tuning

Previously, **glibc** used the per-thread CPU affinity mask for tuning the maximum arena count for **malloc**. As a consequence, restricting the thread affinity mask to a small subset of CPUs in the system could lead to performance degradation.

glibc has been changed to use the configured number of CPUs for determining the maximum arena count. As a result, applications use a larger number of arenas, even when running with a restricted per-thread CPU affinity mask, and the performance degradation no longer occurs.

Jira:RHEL-17157[1]

Improved **glibc** compatibility with applications using **dlclose** on shared objects involved in a dependency cycle

Previously, when unloading a shared object in a dependency cycle using the **dlclose** function in **glibc**, that object's ELF destructor might not have been called before all other objects were unloaded. As a consequence of this late ELF destructor execution, applications experienced crashes and other errors due to the initial shared object's dependencies already being deinitialized.

With this update, **glibc** has been fixed to first call the ELF destructor of the immediate object being unloaded before any other ELF destructors are executed. As a result, compatibility with applications using **dlclose** on shared objects involved in a dependency cycle is improved and crashes no longer occur.

Jira:RHEL-2491[1]

make no longer tries to run directories

Previously, **make** did not check if an executable it was trying to run was actually an executable. Consequently, if the path included a directory with the same name as the executable, **make** tried to run the directory instead. With this update, **make** now does additional checks when searching for an executable. As a result, **make** no longer tries to run directories.

Jira:RHEL-22829

Improved **glibc** wide-character write performance

Previously, the wide **stdio** stream implementation in **glibc** did not treat the default buffer size as large enough for wide-character write operations and used a 16-byte fallback buffer instead, negatively impacting performance. With this update, buffer management is fixed and the entire write buffer is used. As a result, **glibc** wide-character write performance is improved.

Jira:RHEL-19862[1]

The **glibc** getaddrinfo function now correctly reads **ncsd** cache information

Previously, a bug in the **glibc getaddrinfo** function would cause it to occasionally return empty elements in the list address information structure. With this update, the **getaddrinfo** function has been fixed to read and translate **ncsd** cache data correctly and, as a result, returns correct address information.

Jira:RHEL-16643

Improved glibc compatibility with applications using dlclose on shared objects involved in a dependency cycle

Previously, when unloading a shared object in a dependency cycle using the **dlclose** function in **glibc**, that object's ELF destructor might not have been called before all other objects were unloaded. As a consequence of this late ELF destructor execution, applications experienced crashes and other errors due to the initial shared object's dependencies already being deinitialized.

With this update, **glibc** has been fixed to first call the ELF destructor of the immediate object being unloaded before any other ELF destructors are executed. As a result, compatibility with applications using **dlclose** on shared objects involved in a dependency cycle is improved and crashes no longer occur.

Jira:RHEL-12362

ncsd no longer fails to start due to inconsistent cache expiry information

Previously, the **glibc** Name Service Switch Caching Daemon (**nscd**) could fail to start due to inconsistent cache expiry information in the persistent cache file. With this update, **ncsd** now marks cache entries with inconsistent timing information for deletion and skips them. As a result, **ncsd** no longer fails to start due to inconsistent cache expiry information.

Jira:RHEL-3397

Consistently fast glibc thread-local storage performance

Previously, the **glibc** dynamic linker did not adjust certain thread-local storage (TLS) metadata after shared objects with TLS were loaded by using the **dlopen()** function, which consequently caused slow TLS access. With this update, the dynamic linker now updates TLS metadata for TLS changes caused by **dlopen()** calls. As a result, TLS access is consistently fast.

Jira:RHEL-2123

8.12. IDENTITY MANAGEMENT

Allocated memory now released when an operation is completed

Previously, memory allocated by the KCM for each operation was not being released until the connection was closed. As a result, for client applications that opened a connection and ran many operations on the same connection, it led to a noticeable memory increase because the allocated memory was not released until the connection closed. With this update, the memory allocated for an operation is now released as soon as the operation is completed.

Jira:SSSD-7015

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

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

With this update, a case-sensitive comparison is replaced with a case-insensitive comparison that ignores the case of a character. As a result, IdM clients can now lookup users of an AD trusted domain, even if their usernames contain mixed case characters and they are configured with overrides in IdM.

Jira:SSSD-6096

SSSD correctly returns an error if no grace logins remain while changing a password

Previously, if a user's LDAP password had expired, SSSD tried to change the password even after the initial bind of the user failed as there were no more grace logins left. However, the error returned to the user did not indicate the reason for the failure. With this update, the request to change the password is aborted if the bind fails and SSSD returns an error message indicating there are no more grace logins and the password must be changed by another means.

Jira:SSSD-6184

Removing systems from a domain using the realm leave command

Previously, if multiple names were set for the **ad_server** option in the **sssd.conf** file, running the **realm leave** command resulted in parsing errors and the system was not removed from the domain. With this update, the **ad_server** option is properly evaluated and the correct domain controller name is used and the system is correctly removed from the domain.

Jira:SSSD-6081

KCM logs to the correct **sssd.kcm.log** file

Previously, **logrotate** correctly rotated the Kerberos Credential Manager (KCM) log files but KCM incorrectly wrote the logs to the old log file, **sssd_kcm.log.1**. If KCM was restarted, it used the correct log file. With this update, after **logrotate** is invoked, log files are rotated and KCM correctly logs to the **sssd_kcm.log** file.

Jira:SSSD-6652

The **realm leave --remove** command no longer asks for credentials

Previously, the **realm** utility did not correctly check if a valid Kerberos ticket was available when running the **realm leave** operation. As a result, users were asked to enter a password even though a valid Kerberos ticket was available. With this update, **realm** now correctly verifies if there is a valid Kerberos ticket and no longer requests the user to enter a password when running the **realm leave --remove** command.

Jira:SSSD-6425

KDC now runs extra checks when general constrained delegation requests is processed

Previously, the **forwardable** flag in Kerberos tickets issued by KDCs running on Red Hat Enterprise Linux 8 was vulnerable, allowing unauthorized modification without detection. This vulnerability could lead to impersonation attacks, even from/by users without specific privileges. With this update, KDC runs extra checks when it processes general constrained delegation requests, ensuring detection and rejection of unauthorized flag modifications, thus removing the vulnerability.

Jira:RHEL-9984[1]

Check on the forwardable flag is disabled in cases where SIDs are generated for the domain

Previously, the update providing a fix for CVE-2020-17049 relied on the Kerberos PAC to run certain checks on the ticket **forwardable** flag when the KDC processes a general constrained delegation request. However, the PAC is generated only on domains where the SIDs generation task was executed

in the past. While this task is automatically performed for all IdM domains created on Red Hat Enterprise Linux (RHEL) 8.5 and newer, domains initialized on older versions require manual execution of this task.

In case the SIDs generation task was never executed manually for IdM domains initialized on RHEL 8.4 and older, the PAC will be missing on Kerberos tickets, resulting in rejection of all general constrained delegation requests. This includes IdM's HTTP API, which relies on general constrained delegation.

With this update, the check of the **forwardable** flag is disabled in cases where SIDs were not generated for the domain. Services relying on general constrained delegation, including IdM HTTP API, continue working. However, Red Hat recommends running the SIDs generation task on the domain as soon as possible, especially if the domain has custom general constrained delegation rules configured. Until this is done, the domain remains vulnerable to CVE-2020-17049.

Jira:RHEL-22313

IdM Vault encryption and decryption no longer fails in FIPS mode

Previously, IdM Vault used OpenSSL RSA-PKCS1v15 as the default padding wrapping algorithm. However, none of the FIPS certified modules in RHEL supported PKCS#1 v1.5 as a FIPS approved algorithm, causing IdM Vault to fail in FIPS mode. With this update, IdM Vault supports the RSA-OAEP padding wrapping algorithm as a fallback. As a result, IdM Vault encryption and decryption now work correctly in FIPS mode.

Jira:RHEL-12143[1]

Directory Server no longer fails after abandoning the paged result search

Previously, a race condition was a reason for heap corruption and Directory Server failure during abandoning paged result search. With this update, the race condition was fixed, and Directory Server failure no longer occurs.

Jira:RHEL-16830[1]

If the **nsslapd-numlisteners** attribute value is more than 2, Directory Server no longer fails

Previously, if the **nsslapd-numlisteners** attribute value was higher than **2**, Directory Server sometimes closed the listening file descriptor instead of the accepted file descriptor. As a consequence, a segmentation fault occurred in Directory Server. With this update, Directory Server closes the correct descriptor and continues listening on ports correctly.

Jira:RHEL-17175

The autobind operation now does not impacts operations performed on other connections

Previously, when the autobind operation was in progress, Directory Server stopped listening to new operations on any connection. With this update, the autobind operation does not impact the operations performed on the other connection.

Jira:RHEL-5111

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

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

Bugzilla:2094673

8.13. THE WEB CONSOLE

VNC console now works at most resolutions

Previously, when using the Virtual Network Computing (VNC) console under certain display resolutions, a mouse offset problem was present or only a part of the interface was visible. Consequently, using the VNC console was not possible.

With this update, the problem has been fixed and the VNC console works correctly at most resolutions, with the exception of ultra high resolutions, such as 3840 x 2160 px.

Note that a small offset between the recorded and displayed positions of the cursor might still be present. However, this does not significantly impact the usability of the VNC console.

Bugzilla:2030836

8.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Cluster start no longer times out when the SBD delay-start value is high

Previously, when a user configured SBD fencing in a cluster by using the **ha_cluster** System Role and set the **delay-start** option to a value close to or higher than 90 seconds, the cluster start timed out. This is because the default **systemd** start timeout is 90 seconds, which the system reached before the SBD start delay value. With this fix, the **ha_cluster** System Role overrides the **sbd.service** start timeout in **systemd** so that it is higher than the value of **delay-start**. This allows the system to start successfully even with high values of the **delay-start** option.

Jira:RHEL-18026[1]

network role validates routing rules with 0.0.0.0/0 or ::/0

Previously, when the **from:** or **to:** settings were set to the **0.0.0.0/0** or **::/0** addresses in the routing rule, the **network** RHEL System Role failed to configure the routing rule and rejected the settings as invalid. With this update, the **network** role allows **0.0.0.0/0** and **::/0** for **from:** and **to:** in routing rule validation. As a result, the role successfully configures the routing rules without raising the validation errors.

Jira:RHEL-1683

Running read-scale clusters and installing mssql-server-ha no longer requires certain variables

Previously, if you used the **mssql** RHEL System Role to configure a read-scale cluster without certain variables (**mssql_ha_virtual_ip**, **mssql_ha_login**, **mssql_ha_login_password**, and **mssql_ha_cluster_run_role**), the role failed with an error message "Variable not defined". However, these variables are not necessary to run a read-scale cluster. The role also tried to install the **mssql-server-ha**, which is not required for a read-scale cluster. With this fix, the requirement for these variables was removed. As a result, running a read-scale cluster proceeds successfully without the error message.

Jira:RHEL-3540

The Kdump system role works correctly when the kexec_crash_size file is busy

The **/sys/kernel/kexec_crash_size** file provides the size of the memory region allocated for crash kernel memory.

Previously, the Kdump system role failed when the **/sys/kernel/kexec_crash_size** file was busy. With this update, the system role retries reading the file when it is available. As a result, the system role no longer fails when the file is busy.

Jira:RHEL-3353

selinux role no longer uses the item loop variable

Previously, the **selinux** RHEL System Role used the **item** loop variable. This might have resulted in the following warning message when you called the **selinux** role from another role:

> [WARNING]: TASK: fedora.linux_system_roles.selinux : Restore SELinux labels on filesystem tree: The loop variable 'item' is already in use.
> You should set the `loop_var` value in the `loop_control` option for the task to something else to avoid variable collisions and unexpected behavior.

With this release, the **selinux** role uses **__selinux_item** as a loop variable. As a result, the warning that the **item** variable is already in use is no longer displayed even if you call the **selinux** role from another role.

Jira:RHEL-19040

The ha_cluster system role now correctly configures a firewall on a qnetd host

Previously, when a user configured a **qnetd** host and set the **ha_cluster_manage_firewall** variable to **true** by using the **ha_cluster** system role, the role did not enable high-availability services in the firewall. With this fix, the **ha_cluster** system role now correctly configures a firewall on a **qnetd** host.

Jira:RHEL-17875

The postgresql RHEL System Role now installs the correct version of PostgreSQL

Previously, if you tried to run the **postgresql** RHEL System Role with the **postgresql_version: "15"** variable defined on a RHEL managed node, PostgreSQL version 13 was installed instead of version 15. This bug has been fixed, and the **postgresql** role installs the version set in the variable.

Jira:RHEL-5274

keylime_server role correctly reports registrar service status

Previously, when the **keylime_server** role playbook provided incorrect information, the role incorrectly reported the start as successful. With this update, the role now correctly reports a failure when incorrect information is provided, and the timeout when waiting for opened ports has been reduced from approximately 300 seconds to approximately 30 seconds.

Jira:RHEL-15909

The podman RHEL system role now sets and cancels linger properly for rootless containers

Previously, the **podman** RHEL System Role did not set and cancel linger properly for rootless containers. Consequently, deploying secrets or containers for rootless users produced errors in some cases, and failed to cancel linger when removing resources in some cases. With this update, the **podman** RHEL System Role ensures that linger is enabled for rootless users before doing any secret or container resource management, and ensures that linger is canceled for rootless users when there are no more secrets or container resources to be managed. As a result, the role correctly manages lingering for rootless users.

Jira:RHEL-22228

nbde_server role now works with socket overrides

Previously, the **nbde_server** RHEL System Role assumed that the only file in the **tangd** socket override directory was the **override.conf** file for a custom port. Consequently, the role deleted the directory if there was no port customization without checking other files, and the system re-created the directory in subsequent runs.

With this release, the role has been fixed to prevent changing attributes of the port override file and deleting the directory if there are other files. As a result, the role correctly works if **tangd** socket override files are managed also outside of the role.

Jira:RHEL-25508

A volume quadlet service name no longer fails

Previously, starting the volume service name produced an error similar to the following one: "Could not find the requested service NAME.volume: host" With this update, the volume quadlet service name is changed to **basename-volume.service**. As a result, the volume service starts with no errors.

For more information, see Volume unit man page.

Jira:RHEL-21401

Ansible now preserves JSON strings for use in secrets

Previously, Ansible converted JSON strings to the corresponding JSON object if the value was used in a loop and strings similar to **data: "{{ value }}"** As a consequence, you cannot pass JSON strings as secrets and have the value preserved. This update casts the data value to a string when passing to the **podman_secret** module. As a result, JSON strings are preserved as-is for use in secrets.

Jira:RHEL-22309

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

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

Bugzilla:2186218

8.15. VIRTUALIZATION

RT VMs with a FIFO scheduler now boots correctly

Previously, after setting a real-time (RT) virtual machine (VM) to use the **fifo** setting for the vCPU scheduler, the VM became unresponsive when you attempted to boot it. Instead, the VM displayed the **Guest has not initialized the display (yet)** error. With this update, the error has been fixed, and setting **fifo** for the vCPU scheduler works as expected in the described circumstances.

Jira:RHEL-2815[1]

A dump failure no longer blocks IBM Z VMs with Secure Execution from running

Previously, when a dump of an IBM Z virtual machine (VM) with Secure Execution failed, the VM remained in a paused state and was blocked from running. For example, dumping a VM by using the **virsh dump** command fails if there is not enough space on the disk.

The underlying code has been fixed and Secure Execution VMs resume operation successfully after a dump failure.

Jira:RHEL-16695[1]

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

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

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

Bugzilla:1926147[1]

Using a large number of queues no longer causes VMs to fail

Previously, virtual machines (VMs) might have failed when the virtual Trusted Platform Module (vTPM) device was enabled and the multi-queue virtio-net feature was configured to use more than 250 queues.

This problem was caused by a limitation in the vTPM device. With this update, the problem has been fixed and VMs with more than 250 queues and with vTPM enabled now work reliably.

Jira:RHEL-13335[1]

Windows guests boot more reliably after a v2v conversion on hosts with AMD EPYC CPUs

After using the **virt-v2v** utility to convert a virtual machine (VM) that uses Windows 11 or a Windows Server 2022 as the guest OS, the VM previously failed to boot. This occurred on hosts that use AMD EPYC series CPUs. Now, the underlying code has been fixed and VMs boot as expected in the described circumstances.

Bugzilla:2168082[1]

nodedev-dumpxml lists attributes correctly for certain mediated devices

Before this update, the **nodedev-dumpxml** utility did not list attributes correctly for mediated devices that were created using the **nodedev-create** command. This has been fixed, and **nodedev-dumpxml** now displays the attributes of the affected mediated devices properly.

Bugzilla:2143158

virtiofs devices could not be attached after restarting virtqemud or libvirtd

Previously, restarting the **virtqemud** or **libvirtd** services prevented **virtiofs** storage devices from being attached to virtual machines (VMs) on your host. This bug has been fixed, and you can now attach **virtiofs** devices in the described scenario as expected.

Bugzilla:2078693

Hot plugging a Watchdog card to a virtual machine no longer fails

Previously, if no PCI slots were available, adding a Watchdog card to a running virtual machine (VM) failed with the following error:

> Failed to configure watchdog
> ERROR Error attempting device hotplug: internal error: No more available PCI slots

With this update, the problem has been fixed and adding a Watchdog card to a running VM now works as expected.

Bugzilla:2173584

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

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

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

Jira:RHEL-7135

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

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

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

Jira:RHEL-1860[1]

CHAPTER 9. TECHNOLOGY PREVIEWS

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

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

9.1. INSTALLER AND IMAGE CREATION

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

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

Jira:RHEL-10216[1]

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

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

- graphical, text, or cmdline

- ostreecontainer

- clearpart, zerombr

- autopart

- part

- logvol, volgroup

- reboot and shutdown

- lang

- rootpw

- sshkey

- bootloader – Available only with the **--append** optional parameter.

- user

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

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

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

- url (if there is a need to fetch a **stage2** image for installation, for example, PXE installations, use **inst.stage2=** on the kernel instead of providing a url for **stage2** inside the Kickstart file)

- liveimg

- vnc

- authconfig and authselect (provide relevant configuration in the container image instead)

- module

- repo

- zipl

- zfcp

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

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

Jira:RHEL-2250[1]

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

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

Jira:RHEL-17205[1]

The bootc image builder tool is available as a Technology Preview

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

Jira:RHELDOCS-17468[1]

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

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

Jira:RHELDOCS-17733[1]

9.2. SECURITY

gnutls now uses kTLS as a Technology Preview

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

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

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

Bugzilla:2108532[1]

The io_uring interface is available as a Technology Preview

io_uring is a new and effective asynchronous I/O interface, which is now available as a Technology Preview. By default, this feature is disabled. You can enable this interface by setting the **kernel.io_uring_disabled** sysctl variable to any one of the following values:

0

All processes can create **io_uring** instances as usual.

1

io_uring creation is disabled for unprivileged processes. The **io_uring_setup** fails with the **-EPERM** error unless the calling process is privileged by the **CAP_SYS_ADMIN** capability. Existing **io_uring** instances can still be used.

2

io_uring creation is disabled for all processes. The **io_uring_setup** always fails with **-EPERM**. Existing **io_uring** instances can still be used. This is the default setting.

An updated version of the SELinux policy to enable the **mmap** system call on anonymous inodes is also required to use this feature.

By using the **io_uring** command pass-through, an application can issue commands directly to the underlying hardware, such as **nvme**.

Jira:RHEL-11792[1]

9.3. RHEL FOR EDGE

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

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

Jira:RHELDOCS-17752[1]

9.4. SHELLS AND COMMAND-LINE TOOLS

GIMP available as a Technology Preview in RHEL 9

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

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

Bugzilla:2047161[1]

9.5. INFRASTRUCTURE SERVICES

Socket API for TuneD available as a Technology Preview

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

Bugzilla:2113900

9.6. NETWORKING

WireGuard VPN is available as a Technology Preview

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

For further details, see Setting up a WireGuard VPN .

Bugzilla:1613522[1]

kTLS available as a Technology Preview

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

Bugzilla:1570255[1]

The **systemd-resolved** service is available as a Technology Preview

The **systemd-resolved** service provides name resolution to local applications. The service implements a caching and validating DNS stub resolver, a Link-Local Multicast Name Resolution (LLMNR), and Multicast DNS resolver and responder.

Note that **systemd-resolved** is an unsupported Technology Preview.

Bugzilla:2020529

The PRP and HSR protocols are now available as a Technology Preview

This update adds the **hsr** kernel module that provides the following protocols:

- Parallel Redundancy Protocol (PRP)

- High-availability Seamless Redundancy (HSR)

The IEC 62439-3 standard defines these protocols, and you can use this feature to configure zero-loss redundancy in Ethernet networks.

Bugzilla:2177256[1]

NetworkManager and the Nmstate API support MACsec hardware offload

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

Note that this feature is an unsupported Technology Preview.

Jira:RHEL-24337

NetworkManager enables configuring HSR and PRP interfaces

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

Jira:RHEL-5852

Offloading IPsec encapsulation to a NIC is now available as a Technology Preview

This update adds the IPsec packet offloading capabilities to the kernel. Previously, it was possible to only offload the encryption to a network interface controller (NIC). With this enhancement, the kernel can now offload the entire IPsec encapsulation process to a NIC to reduce the workload.

Note that offloading the IPsec encapsulation process to a NIC also reduces the ability of the kernel to monitor and filter such packets.

Bugzilla:2178699[1]

Network drivers for modems in RHEL are available as Technology Preview

Device manufacturers support Federal Communications Commission (FCC) locking as the default setting. FCC provides a lock to bind WWAN drivers to a specific system where WWAN drivers provide a channel to communicate with modems. Based on the modem PCI ID, manufacturers integrate unlocking tools on Red Hat Enterprise Linux for ModemManager. However, a modem remains unusable if not unlocked previously even if the WWAN driver is compatible and functional. Red Hat Enterprise Linux provides the drivers for the following modems with limited functionality as a Technology Preview:

- Qualcomm MHI WWAM MBIM - Telit FN990Axx

- Intel IPC over Shared Memory (IOSM) - Intel XMM 7360 LTE Advanced

- Mediatek t7xx (WWAN) - Fibocom FM350GL

- Intel IPC over Shared Memory (IOSM) - Fibocom L860GL modem

Jira:RHELDOCS-16760[1], Jira:RHEL-6564, Bugzilla:2110561, Bugzilla:2123542, Bugzilla:2222914

Segment Routing over IPv6 (SRv6) is available as a Technology Preview

The RHEL kernel provides Segment Routing over IPv6 (SRv6) as a Technology Preview. You can use this functionality to optimize traffic flows in edge computing or to improve network programmability in data centers. However, the most significant use case is the end-to-end (E2E) network slicing in 5G deployment scenarios. In that area, the SRv6 protocol provides you with the programmable custom network slices and resource reservations to address network requirements for specific applications or services. At the same time, the solution can be deployed on a single-purpose appliance, and it satisfies the need for a smaller computational footprint.

Bugzilla:2186375[1]

kTLS rebased to version 6.3

The kernel Transport Layer Security (KTLS) functionality is a Technology Preview. In RHEL 9.3, kTLS was rebased to the 6.3 upstream version, and notable changes include:

- Added the support for 256-bit keys with TX device offload

- Delivered various bug fixes

Bugzilla:2183538[1]

9.7. KERNEL

The Soft-iWARP driver is available as a Technology Preview

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

Bugzilla:2023416[1]

SGX available as a Technology Preview

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

- Modifying EPCM permissions of regular enclave pages that belong to an initialized enclave.

- Dynamic addition of regular enclave pages to an initialized enclave.

- Expanding an initialized enclave to accommodate more threads.

- Removing regular and TCS pages from an initialized enclave.

Bugzilla:1660337[1]

rvu_af, rvu_nicpf, and rvu_nicvf available as Technology Preview

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

rvu_nicpf

Marvell OcteonTX2 NIC Physical Function driver

rvu_nicvf

Marvell OcteonTX2 NIC Virtual Function driver

rvu_nicvf

Marvell OcteonTX2 RVU Admin Function driver

Bugzilla:2040643[1]

python-drgn available as a Technology Preview

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

Jira:RHEL-6973[1]

The IAA crypto driver is now available as a Technology Preview

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

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

For details about the IAA crypto driver, see:

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

- IAA Compression Accelerator Crypto Driver

Jira:RHEL-20145[1]

9.8. FILE SYSTEMS AND STORAGE

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

In RHEL 9, the DAX file system is available as a Technology Preview. DAX provides means for an application to directly map persistent memory into its address space. To use DAX, a system must have some form of persistent memory available, usually in the form of one or more Non-Volatile Dual In-line Memory Modules (NVDIMMs), and a DAX compatible file system must be created on the NVDIMM(s). Also, the file system must be mounted with the **dax** mount option. Then, an **mmap** of a file on the dax-mounted file system results in a direct mapping of storage into the application's address space.

Bugzilla:1995338[1]

NVMe-oF Discovery Service features available as a Technology Preview

The NVMe-oF Discovery Service features, defined in the NVMexpress.org Technical Proposals (TP)

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

Bugzilla:2021672[1]

nvme-stas package available as a Technology Preview

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

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

Bugzilla:1893841[1]

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

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

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

Bugzilla:2027304[1]

9.9. COMPILERS AND DEVELOPMENT TOOLS

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

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

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

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

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

Bugzilla:1980981

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

With this update, when comparing binaries, you can suppress warnings related to fake flexible arrays that were converted to true flexible arrays by using the following suppression specification:

- - - - type_kind = struct has_size_change = true
has_strict_flexible_array_data_member_conversion = true

- - - -

Jira:RHEL-16629[1]

9.10. IDENTITY MANAGEMENT

DNSSEC available as Technology Preview in IdM

Identity Management (IdM) servers with integrated DNS now implement DNS Security Extensions (DNSSEC), a set of extensions to DNS that enhance security of the DNS protocol. DNS zones hosted on IdM servers can be automatically signed using DNSSEC. The cryptographic keys are automatically generated and rotated.

Users who decide to secure their DNS zones with DNSSEC are advised to read and follow these documents:

- DNSSEC Operational Practices, Version 2

- Secure Domain Name System (DNS) Deployment Guide

- DNSSEC Key Rollover Timing Considerations

Note that IdM servers with integrated DNS use DNSSEC to validate DNS answers obtained from other DNS servers. This might affect the availability of DNS zones that are not configured in accordance with recommended naming practices.

Bugzilla:2084180

ACME available as a Technology Preview

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

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

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

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

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

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

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

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

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

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

Bugzilla:2084181[1]

9.11. DESKTOP

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

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

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

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

- The Firefox web browser

- Red Hat Subscription Manager (**subscription-manager-cockpit**)

- Firewall Configuration (**firewall-config**)

- Disk Usage Analyzer (**baobab**)

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

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

Jira:RHELPLAN-27394[1]

GNOME for the IBM Z architecture available as a Technology Preview

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

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

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

- The Firefox web browser

- Red Hat Subscription Manager (**subscription-manager-cockpit**)

- Firewall Configuration (**firewall-config**)

- Disk Usage Analyzer (**baobab**)

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

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

Jira:RHELPLAN-27737[1]

9.12. THE WEB CONSOLE

The RHEL web console can now manage WireGuard connections

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

Jira:RHELDOCS-17520[1]

9.13. VIRTUALIZATION

Creating nested virtual machines

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

Jira:RHELDOCS-17040[1]

AMD SEV and SEV-ES for KVM virtual machines

As a Technology Preview, RHEL 9 provides the Secure Encrypted Virtualization (SEV) feature for AMD EPYC host machines that use the KVM hypervisor. If enabled on a virtual machine (VM), SEV encrypts the VM's memory to protect the VM from access by the host. This increases the security of the VM.

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

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

Jira:RHELPLAN-65217[1]

Intel TDX in RHEL guests

As a Technology Preview, the Intel Trust Domain Extension (TDX) feature can now be used in RHEL 9.2 and later guest operating systems. If the host system supports TDX, you can deploy hardware-isolated RHEL 9 virtual machines (VMs), called trust domains (TDs). Note, however, that TDX currently does not work with **kdump**, and enabling TDX will cause **kdump** to fail on the VM.

Bugzilla:1955275[1]

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

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

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

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

Bugzilla:2142102[1]

Intel vGPU available as a Technology Preview

As a Technology Preview, it is possible to divide a physical Intel GPU device into multiple virtual devices referred to as **mediated devices**. These mediated devices can then be assigned to multiple virtual machines (VMs) as virtual GPUs. As a result, these VMs share the performance of a single physical Intel GPU.

Note that this feature is deprecated and was removed entirely with the RHEL 9.3 release.

Jira:RHELDOCS-17050[1]

CPU clusters on ARM 64

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

Jira:RHEL-7043[1]

Live migrating a VM with a Mellanox virtual function is now available as a Technology Preview

As a Technology Preview, you can now live migrate a virtual machine (VM) with an attached virtual function (VF) of a Mellanox networking device.

This feature is currently available only on a Mellanox CX-7 networking device. The VF on the Mellanox CX-7 networking device uses a new **mlx5_vfio_pci** driver, which adds functionality that is necessary for the live migration, and **libvirt** binds the new driver to the VF automatically.

Jira:RHEL-13007[1]

9.14. RHEL IN CLOUD ENVIRONMENTS

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

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

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

Jira:RHELPLAN-139800[1]

9.15. CONTAINERS

The **podman-machine** command is unsupported

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

Jira:RHELDOCS-16861[1]

Building multi-architecture images is available as a Technology Preview

The **podman farm build** command, which you can use to create multi-architecture container images, is available as a Technology Preview.

A farm is a group of machines that have a unix podman socket running in them. The nodes in the farm can have different machines of different architectures. The **podman farm build** command is faster than the **podman build --arch --platform** command.

You can use **podman farm build** to perform the following actions:

- Build an image on all nodes in a farm.

- Bundle nodes up into a manifest list.

- Execute the **podman build** command on all the farm nodes.

- Push the images to the registry specified by using the **--tag** option.

- Locally create a manifest list.

- Push the manifest list to the registry.
  The manifest list contains one image per native architecture type that is present in the farm.

Jira:RHELPLAN-154436[1]

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

The **rhel9/rhel-bootc** container image is now available in the Red Hat Container Registry as a Technology Preview. With the RHEL bootable container images, you can build, test, and deploy an operating system exactly as a container. The RHEL bootable container images differ from the existing application Universal Base Images (UBI) thanks to the following enhancements: RHEL bootable container images contain additional components necessary to boot, such as, kernel, initrd, bootloader, firmware, between others. There are no changes to existing container images. For more information, see Red Hat Ecosystem Catalog .

Jira:RHELDOCS-17803[1]

The composefs filesystem is now available as a Technology Preview

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

Jira:RHEL-18157[1]

CHAPTER 10. DEPRECATED FUNCTIONALITY

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

Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments. For the most recent list of deprecated functionality within a particular major release, refer to the latest version of release documentation.

The support status of deprecated functionality remains unchanged within Red Hat Enterprise Linux 9. For information about the length of support, see Red Hat Enterprise Linux Life Cycle and Red Hat Enterprise Linux Application Streams Life Cycle.

Deprecated hardware components are not recommended for new deployments on the current or future major releases. Hardware driver updates are limited to security and critical fixes only. Red Hat recommends replacing this hardware as soon as reasonably feasible.

A package can be deprecated and not recommended for further use. Under certain circumstances, a package can be removed from a product. Product documentation then identifies more recent packages that offer functionality similar, identical, or more advanced to the one deprecated, and provides further recommendations.

For information regarding functionality that is present in RHEL 8 but has been *removed* in RHEL 9, see Considerations in adopting RHEL 9 .

10.1. INSTALLER AND IMAGE CREATION

Deprecated Kickstart commands

The following Kickstart commands have been deprecated:

- **timezone --ntpservers**

- **timezone --nontp**

- **logging --level**

- **%packages --excludeWeakdeps**

- **%packages --instLangs**

- **%anaconda**

- **pwpolicy**

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

Bugzilla:1899167[1]

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

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

Note that specifying a user or group customization in blueprints that are used to deploy an existing OSTree commit, such as **edge-raw-image**, **edge-installer**, and **edge-simplified-installer** image types remains supported.

Bugzilla:2173928

The initial-setup package now has been deprecated

The **initial-setup** package has been deprecated in Red Hat Enterprise Linux 9.3 and will be removed in the next major RHEL release. As a replacement, use **gnome-initial-setup** for the graphical user interface.

Jira:RHELDOCS-16393[1]

The provider_hostip and provider_fedora_geoip values of the inst.geoloc boot option are deprecated

The **provider_hostip** and **provider_fedora_geoip** values that specified the GeoIP API for the **inst.geoloc=** boot option are deprecated. As a replacement, you can use the **geolocation_provider=URL** option to set the required geolocation in the installation program configuration file. You can still use the **inst.geoloc=0** option to disable the geolocation.

Bugzilla:2127473

Capturing screenshots from the Anaconda GUI with a global hotkey is deprecated

Previously, users could capture screenshots of the Anaconda GUI by using a global hotkey. This meant that users could extract the screenshots manually from the installation environment for any further usage. This functionality has been deprecated.

Jira:RHELDOCS-17166[1]

Anaconda built-in help has been deprecated

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

Jira:RHELDOCS-17309[1]

Support for NVDIMM devices has been deprecated

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

Jira:RHELDOCS-17702

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

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

As a workaround, use the **modprobe.blacklist=** kernel command line option together with the **inst.dd**

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

Jira:RHEL-4762

10.2. SECURITY

SHA-1 is deprecated for cryptographic purposes

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

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

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

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

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

Jira:RHELPLAN-110763[1]

fapolicyd.rules is deprecated

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

Bugzilla:2054740

SCP is deprecated in RHEL 9

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

- In the **scp** utility, SCP is replaced by the SSH File Transfer Protocol (SFTP) by default.

- The OpenSSH suite does not use SCP in RHEL 9.

- SCP is deprecated in the **libssh** library.

Jira:RHELPLAN-99136[1]

OpenSSL requires padding for RSA encryption in FIPS mode

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

Bugzilla:2168665

OpenSSL deprecates the Engines API

The OpenSSL 3.0 TLS toolkit deprecated the Engines API. The Engines interface is superseded by the Providers API. The migration of applications and existing engines to Providers is underway. The deprecated Engines API may be removed in a future major release.

Jira:RHELDOCS-17958[1]

openssl-pkcs11 is now deprecated

As a part of the ongoing migration of deprecated OpenSSL engines to the Providers API, the **pkcs11-provider** package replaces the **openssl-pkcs11** package (**engine_pkcs11**). The **openssl-pkcs11** package is now deprecated. The **openssl-pkcs11** package may be removed in a future major release.

Jira:RHELDOCS-16716[1]

RHEL 8 and 9 OpenSSL certificate and signing containers are now deprecated

The OpenSSL portable certificate and signing containers available in the **ubi8/openssl** and **ubi9/openssl** repositories in the Red Hat Ecosystem Catalog are now deprecated due to low demand.

Jira:RHELDOCS-17974[1]

Digest–MD5 in SASL is deprecated

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

Bugzilla:1995600[1]

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

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

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

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

Bugzilla:1975836

/etc/system-fips is now deprecated

Support for indicating FIPS mode through the **/etc/system-fips** file has been removed, and the file will not be included in future versions of RHEL. To install RHEL in FIPS mode, add the **fips=1** parameter to

the kernel command line during the system installation. You can check whether RHEL operates in FIPS mode by using the **fips-mode-setup --check** command.

Jira:RHELPLAN-103232[1]

libcrypt.so.1 is now deprecated

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

Bugzilla:2034569

10.3. SUBSCRIPTION MANAGEMENT

The deprecated **--token** option of **subscription-manager register** will stop working at the end of November 2024

The deprecated **--token=<TOKEN>** option of the **subscription-manager register** command will no longer be a supported authentication method from the end of November 2024. The default entitlement server, **subscription.rhsm.redhat.com**, will no longer be allowing token-based authentication. As a consequence, if you use **subscription-manager register --token=<TOKEN>**, the registration will fail with the following error message:

> Token authentication not supported by the entitlement server

To register your system, use other supported authorization methods, such as including paired options **--username** / **--password** OR **--org** / **--activationkey** with the **subscription-manager register** command.

Bugzilla:2163716

10.4. SHELLS AND COMMAND-LINE TOOLS

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

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

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

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

Bugzilla:1997366[1]

The SQLite database backend in Bacula has been deprecated

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

Jira:RHEL-6856

The **%vmeff** metric from the **sysstat** package has been deprecated

The **%vmeff** metric from the **sysstat** package to measure the page reclaim efficiency will no longer be supported in a future major version of RHEL. The values of the **%vmeff** column returned by the **sar -B** command are incorrect because **sysstat** does not parse all relevant **/proc/vmstat** values provided by later kernel versions.

You can calculate the **%vmeff** value manually from the **/proc/vmstat** file. For details, see Why the **sar(1)** tool reports **%vmeff** values beyond 100 % in RHEL 8 and RHEL 9?

Jira:RHELDOCS-17015[1]

Setting the **TMPDIR** variable in the ReaR configuration file is deprecated

Setting the **TMPDIR** environment variable in the **/etc/rear/local.conf** or **/etc/rear/site.conf** ReaR configuration file), by using a statement such as **export TMPDIR=…**, is deprecated.

To specify a custom directory for ReaR temporary files, export the variable in the shell environment before executing ReaR. For example, execute the **export TMPDIR=…** statement and then execute the **rear** command in the same shell session or script.

Jira:RHELDOCS-18049[1]

cgroupsv1 is now deprecated in RHEL 9

The **cgroups** is a kernel subsystem used for process tracking, system resource allocation and partitioning. Systemd service manager supports booting in the cgroups **v1** mode as well as in cgroups **v2** mode. In Red Hat Enterprise Linux 9, the default mode is **v2**. In Red Hat Enterprise Linux 10, systemd will not support booting in the cgroups **v1** mode and only cgroups **v2** mode will be available.

Jira:RHELDOCS-17545[1]

10.5. INFRASTRUCTURE SERVICES

Client-side and server-side DHCP packages are deprecated

Internet Systems Consortium (ISC) has announced the end of maintenance for ISC DHCP as of the end of 2022. As a result, Red Hat has decided to deprecate the use of client-side and server-side DHCP packages in RHEL 9 and not to distribute them in later major versions of RHEL. Customers must prepare for the transition to available alternatives, such as **dhcpcd** and **ISC Kea**.

Jira:RHELDOCS-17135[1]

The **sendmail, libotr, mod_security,** and **spamassassin** packages are now deprecated

The following packages are deprecated in RHEL 9 and will not be distributed in later major versions of RHEL:

- sendmail – Red Hat recommends migrating to the postfix mail daemon, which is supported.

- libotr

- mod_security

- spamassassin

Jira:RHEL-22385[1]

10.6. NETWORKING

Network teams are deprecated in RHEL 9

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

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

For details about how to migrate a team to a bond, see Migrating a network team configuration to network bond.

Bugzilla:1935544[1]

NetworkManager connection profiles in ifcfg format are deprecated

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

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

Bugzilla:1894877[1]

The iptables back end in firewalld is deprecated

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

Bugzilla:2089200

The firewalld lockdown feature is deprecated.

The lockdown feature in **firewalld** is deprecated because it cannot prevent processes that are running as **root** from adding themselves to the allow list. The lockdown feature may be removed in a future major RHEL release.

Jira:RHEL-17708

The connection.master, connection.slave-type, and connection.autoconnect-slaves properties are deprecated

Red Hat is committed to using conscious language. For details about this initiative, see Making open source more inclusive. Therefore, the **connection.master**, **connection.slave-type**, and **connection.autoconnect-slaves** properties were renamed. To ensure backward compatibility, aliases have been created that map the old property names to the new ones:

- **connection.master** is an alias for **connection.controller**

- **connection.slave-type** is an alias for **connection.port-type**

- **connection.autoconnect-slaves** is an alias for **connection.autoconnect-ports**

Note that the **connection.master**, **connection.slave-type**, and **connection.autoconnect-slaves** aliases are deprecated and will be removed in a future RHEL version.

Jira:RHEL-17619[1]

The **PF_KEYv2** kernel API is deprecated

Applications can configure the kernel's IPsec implementation by using the **PV_KEYv2** and the newer **netlink** API. **PV_KEYv2** is not actively maintained upstream and misses important security features, such as modern ciphers, offload, and extended sequence number support. As a result, starting with RHEL 9.3, the **PV_KEYv2** API is deprecated and will be removed in the next major RHEL release. If you use this kernel API in your application, migrate it to use the modern **netlink** API as an alternative.

Jira:RHEL-1015[1]

10.7. KERNEL

ATM encapsulation is deprecated in RHEL 9

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

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

Bugzilla:2058153

The **kexec_load** system call for **kexec-tools** has been deprecated

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

For more information, see Is kexec_load supported in RHEL9? .

Bugzilla:2113873[1]

Network teams are deprecated in RHEL 9

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

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

For details about how to migrate a team to a bond, see Migrating a network team configuration to network bond.

Bugzilla:2013884[1]

10.8. FILE SYSTEMS AND STORAGE

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

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

Bugzilla:2038183

Persistent Memory Development Kit (pmdk) and support library have been deprecated in RHEL 9

pmdk is a collection of libraries and tools for System Administrators and Application Developers to simplify managing and accessing persistent memory devices. **pmdk** and support library have been deprecated in RHEL 9. This also includes the **-debuginfo** packages.

The following list of binary packages produced by **pmdk**, including the **nvml** source package have been deprecated:

- **libpmem**
- **libpmem-devel**
- **libpmem-debug**
- **libpmem2**
- **libpmem2-devel**
- **libpmem2-debug**
- **libpmemblk**
- **libpmemblk-devel**
- **libpmemblk-debug**
- **libpmemlog**
- **libpmemlog-devel**
- **libpmemlog-debug**
- **libpmemobj**
- **libpmemobj-devel**
- **libpmemobj-debug**
- **libpmempool**
- **libpmempool-devel**
- **libpmempool-debug**
- **pmempool**
- **daxio**

- **pmreorder**

- **pmdk-convert**

- **libpmemobj++**

- **libpmemobj++-devel**

- **libpmemobj++-doc**

Jira:RHELDOCS-16432[1]

The **md-linear** and **md-faulty** modules have been deprecated

The following MD RAID kernel modules have been deprecated, and will be removed in a future major RHEL release:

- **CONFIG_MD_LINEAR** or **md-linear** module to concatenate multiple drives so that when a single member disk becomes full, data will be written to the next disk until all disks are full.

- **CONFIG_MD_FAULTY** or **md-faulty** module to test a block device that occasionally returns read or write errors. It is useful for testing.

Jira:RHEL-30730[1]

The VDO **sysfs** parameters have been deprecated

The Virtual Data Optimizer (VDO) **sysfs** parameters have been deprecated and will be removed in a future major RHEL release. Except for **log_level**, all module-level **sysfs** parameters for the **kvdo** module will be removed. For individual **dm-vdo** targets, all **sysfs** parameters specific to VDO will also be removed. There is no change for the parameters that are common to all DM targets. Configuration values for **dm-vdo** targets, which are currently set by updating the removed module-level parameters, can no longer be changed.

Statistics and configuration values for **dm-vdo** targets will no longer be accessible through **sysfs**. But these values are still accessible by using **dmsetup message stats**, **dmsetup status**, and **dmsetup table** dmsetup commands

Jira:RHEL-30525

10.9. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

libdb has been deprecated

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

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

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

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

Bugzilla:1927780[1], Jira:RHELPLAN-80695, Bugzilla:1974657

10.10. COMPILERS AND DEVELOPMENT TOOLS

Smaller size of keys than 2048 are deprecated by **openssl** 3.0 in Go's FIPS mode

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

Bugzilla:2111072

Some PKCS1 v1.5 modes are now deprecated in Go's FIPS mode

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

Bugzilla:2092016[1]

32-bit packages are deprecated

Linking against 32-bit multilib packages is deprecated. The ***.i686** packages will remain supported for the life cycle of Red Hat Enterprise Linux 9, but will be removed in the next major version of RHEL.

Jira:RHELDOCS-17917[1]

10.11. IDENTITY MANAGEMENT

SHA-1 in OpenDNSSec is now deprecated

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

Bugzilla:1979521

The SSSD implicit files provider domain is disabled by default

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

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

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

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

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

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

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

2. Configure the name services switch.

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

Jira:RHELPLAN-100639[1]

The SSSD **files** provider has been deprecated

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

Jira:RHELPLAN-139805[1]

The **enumeration** feature has been deprecated for AD and IdM

The **enumeration** feature enables you to list all users or groups by using **getent passwd** or **getent group** commands without arguments for Active Directory (AD), Identity Management (IdM), and LDAP providers. Support for the **enumeration** feature has been deprecated for AD and IdM in Red Hat Enterprise Linux (RHEL) 9. The **enumeration** feature will be removed for AD and IdM in RHEL 10.

Jira:SSSD-6596

The **libsss_simpleifp** subpackage has been deprecated

The **libsss_simpleifp** subpackage that provides the **libsss_simpleifp.so** library has been deprecated in Red Hat Enterprise Linux (RHEL) 9. The **libsss_simpleifp** subpackage might be removed from a future release of RHEL.

Jira:SSSD-6601

The SMB1 protocol is deprecated in Samba

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

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

Jira:RHELDOCS-16612[1]

10.12. DESKTOP

GTK 2 is now deprecated

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

- **adwaita-gtk2-theme**

- **gnome-common**

- **gtk2**

- **gtk2-immodules**

- **hexchat**

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

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

Jira:RHELPLAN-131882[1]

LibreOffice is deprecated

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

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

- The official Flatpak package in the Flathub repository: https://flathub.org/apps/org.libreoffice.LibreOffice.

- The official RPM packages: https://www.libreoffice.org/download/download-libreoffice/.

Jira:RHELDOCS-16300[1]

10.13. GRAPHICS INFRASTRUCTURES

Motif has been deprecated

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

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

- **motif**

- **openmotif**

- **openmotif21**

- **openmotif22**

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

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

Jira:RHELPLAN-98983[1]

10.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The `mssql_ha_cluster_run_role` has been deprecated

The **mssql_ha_cluster_run_role** variable has been deprecated. Instead, use the **mssql_manage_ha_cluster** variable.

Jira:RHEL-19092

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

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

Bugzilla:1999770

10.15. VIRTUALIZATION

SecureBoot image verification using SHA1-based signatures is deprecated

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

Bugzilla:1935497[1]

The virtual floppy driver has become deprecated

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

Bugzilla:1965079

qcow2-v2 image format is deprecated

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

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

Bugzilla:1951814

virt-manager has been deprecated

The Virtual Machine Manager application, also known as **virt-manager**, has been deprecated. The RHEL web console, also known as **Cockpit**, is intended to become its replacement in a subsequent release. It is, therefore, recommended that you use the web console for managing virtualization in a GUI. Note, however, that some features available in **virt-manager** might not be yet available in the RHEL web console.

Jira:RHELPLAN-10304[1]

libvirtd has become deprecated

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

Jira:RHELPLAN-113995[1]

Legacy CPU models are now deprecated

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

- For Intel: models before Intel Xeon 55xx and 75xx Processor families (also known as Nehalem)

- For AMD: models before AMD Opteron G4

- For IBM Z: models before IBM z14

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

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

Bugzilla:2060839

RDMA-based live migration is deprecated

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

Jira:RHELPLAN-153267[1]

The Intel vGPU feature has been removed

Previously, as a Technology Preview, it was possible to divide a physical Intel GPU device into multiple virtual devices referred to as **mediated devices**. These mediated devices could then be assigned to multiple virtual machines (VMs) as virtual GPUs. As a result, these VMs shared the performance of a single physical Intel GPU, however only selected Intel GPUs were compatible with this feature.

Since RHEL 9.3, the Intel vGPU feature has been removed entirely.

Bugzilla:2206599[1]

pmem device passthrough has become deprecated

With this update, the non-volatile memory library (**nvml**) packages have become deprecated, and will be removed in a future major version of RHEL. As a consequence, when the package removal occurs, it will no longer be possible to pass persistent memory (**pmem**) devices to the virtual machines (VMs). Note that emulated NVDIMM devices backed by volatile memory or files will still be available, but will not be possible to configure as persistent.

Jira:RHELDOCS-17989

Using Windows Server 2012 or Windows 8 as a guest operating system is not supported

Because Microsoft ended support for the following versions of Windows, Red Hat also removed support for using these versions as a guest operating system in this update.

- Windows 8

- Windows 8.1

- Windows Server 2012

- Windows Server 2012 R2

Jira:RHEL-11810

10.16. CONTAINERS

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

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

For more information, see Red Hat Enterprise Linux Container Compatibility Matrix .

Jira:RHELPLAN-100087[1]

SHA1 hash algorithm within Podman has been deprecated

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

Bugzilla:2069279[1]

rhel9/pause has been deprecated

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

Bugzilla:2106816

The CNI network stack has been deprecated

The Container Network Interface (CNI) network stack is deprecated and will be removed from Podman in a future minor release of RHEL. Previously, containers connected to the single Container Network Interface (CNI) plugin only via DNS. Podman v.4.0 introduced a new Netavark network stack. You can use the Netavark network stack with Podman and other Open Container Initiative (OCI) container management applications. The Netavark network stack for Podman is also compatible with advanced Docker functionalities. Containers in multiple networks can access containers on any of those networks.

For more information, see Switching the network stack from CNI to Netavark .

Jira:RHELDOCS-16756[1]

The Inkscape and LibreOffice Flatpak images are deprecated

The **rhel9/inkscape-flatpak** and **rhel9/libreoffice-flatpak** Flatpak images, which are available as Technology Previews, have been deprecated.

Red Hat recommends the following alternatives to these images:

- To replace **rhel9/inkscape-flatpak**, use the **inkscape** RPM package.

- To replace **rhel9/libreoffice-flatpak**, see the LibreOffice deprecation release note.

Jira:RHELDOCS-17102[1]

pasta as a network name has been deprecated

The support for **pasta** as a network name value is deprecated and will not be accepted in the next major

release of Podman, version 5.0. You can use the **pasta** network name value to create a unique network mode within Podman by employing the **podman run --network** and **podman create --network** commands.

Jira:RHELDOCS-17038[1]

The BoltDB database backend has been deprecated

The BoltDB database backend is deprecated as of RHEL 9.4. In a future version of RHEL, the BoltDB database backend will be removed and will no longer be available to Podman. For Podman, use the SQLite database backend, which is now the default as of RHEL 9.4.

Jira:RHELDOCS-17495[1]

The CNI network stack has been deprecated

The Container Network Interface (CNI) network stack is deprecated and will be removed in a future release. Use the Netavark network stack instead. For more information, see Switching the network stack from CNI to Netavark.

Jira:RHELDOCS-17518[1]

The Podman v5.0 upcoming deprecations

The following will be deprecated in the upcoming Podman v5.0, which will be released in RHEL 9.5 and RHEL 10.0 Beta:

- The BoltDB database backend will be deprecated. The new SQLite database backend is available.

- The **containers.conf** file will be read-only. The system connections and farm information will be stored in the **podman.connections.json** file, managed only by Podman. Podman continues to support the old configuration options such as **[engine.service_destinations]** and the **[farms]** section. You can still add connections or farms manually if needed, however, it is not possible to delete a connection from the **containers.conf** file with the **podman system connection rm** command.

The following changes are planned for RHEL 10.0 Beta:

- The **pasta** network mode will be the default network mode for rootless containers. The **slirp4netns** network mode will be deprecated.

- The cgroupv1 will be deprecated.

- The CNI network stack will be deprecated.

Jira:RHELDOCS-17462[1]

The **rhel9/openssl** has been deprecated

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

Jira:RHELDOCS-18106[1]

10.17. DEPRECATED PACKAGES

This section lists packages that have been deprecated and will probably not be included in a future major release of Red Hat Enterprise Linux.

For changes to packages between RHEL 8 and RHEL 9, see Changes to packages in the *Considerations in adopting RHEL 9* document.

> **IMPORTANT**
>
> The support status of deprecated packages remains unchanged within RHEL 9. For more information about the length of support, see Red Hat Enterprise Linux Life Cycle and Red Hat Enterprise Linux Application Streams Life Cycle .

The following packages have been deprecated in RHEL 9:

- aacraid

- adwaita-gtk2-theme

- af_key

- anaconda-user-help

- autocorr-af

- autocorr-bg

- autocorr-ca

- autocorr-cs

- autocorr-da

- autocorr-de

- autocorr-dsb

- autocorr-el

- autocorr-en

- autocorr-es

- autocorr-fa

- autocorr-fi

- autocorr-fr

- autocorr-ga

- autocorr-hr

- autocorr-hsb

- autocorr-hu

- autocorr-is

- autocorr-it
- autocorr-ja
- autocorr-ko
- autocorr-lb
- autocorr-lt
- autocorr-mn
- autocorr-nl
- autocorr-pl
- autocorr-pt
- autocorr-ro
- autocorr-ru
- autocorr-sk
- autocorr-sl
- autocorr-sr
- autocorr-sv
- autocorr-tr
- autocorr-vi
- autocorr-vro
- autocorr-zh
- cheese
- cheese-libs
- clutter
- clutter-gst3
- clutter-gtk
- cogl
- daxio
- dbus-glib
- dbus-glib-devel
- dhcp-client

- dhcp-common

- dhcp-relay

- dhcp-server

- enchant

- enchant-devel

- eog

- evolution

- evolution-bogofilter

- evolution-devel

- evolution-help

- evolution-langpacks

- evolution-mapi

- evolution-mapi-langpacks

- evolution-pst

- evolution-spamassassin

- festival

- festival-data

- festvox-slt-arctic-hts

- flite

- flite-devel

- firewire-core

- gedit

- gedit-plugin-bookmarks

- gedit-plugin-bracketcompletion

- gedit-plugin-codecomment

- gedit-plugin-colorpicker

- gedit-plugin-colorschemer

- gedit-plugin-commander

- gedit-plugin-drawspaces

- gedit-plugin-findinfiles

- gedit-plugin-joinlines

- gedit-plugin-multiedit

- gedit-plugin-sessionsaver

- gedit-plugin-smartspaces

- gedit-plugin-synctex

- gedit-plugin-terminal

- gedit-plugin-textsize

- gedit-plugin-translate

- gedit-plugin-wordcompletion

- gedit-plugins

- gedit-plugins-data

- ghostscript-x11

- gnome-common

- gnome-photos

- gnome-photos-tests

- gnome-screenshot

- gnome-themes-extra

- gtk2

- gtk2-devel

- gtk2-devel-docs

- gtk2-immodule-xim

- gtk2-immodules

- highcontrast-icon-theme

- inkscape

- inkscape-docs

- inkscape-view

- iptables-devel

- iptables-libs

- iptables-nft

- iptables-nft-services

- iptables-utils

- libdb

- libgdata

- libgdata-devel

- libpmem

- libpmem-debug

- libpmem-devel

- libpmem2

- libpmem2-debug

- libpmem2-devel

- libpmemblk

- libpmemblk-debug

- libpmemblk-devel

- libpmemlog

- libpmemlog-debug

- libpmemlog-devel

- libpmemobj

- libpmemobj-debug

- libpmemobj-devel

- libpmempool

- libpmempool-debug

- libpmempool-devel

- libreoffice

- libreoffice-base

- libreoffice-calc

- libreoffice-core

- libreoffice-data

- libreoffice-draw

- libreoffice-emailmerge

- libreoffice-filters

- libreoffice-gdb-debug-support

- libreoffice-graphicfilter

- libreoffice-gtk3

- libreoffice-help-ar

- libreoffice-help-bg

- libreoffice-help-bn

- libreoffice-help-ca

- libreoffice-help-cs

- libreoffice-help-da

- libreoffice-help-de

- libreoffice-help-dz

- libreoffice-help-el

- libreoffice-help-en

- libreoffice-help-eo

- libreoffice-help-es

- libreoffice-help-et

- libreoffice-help-eu

- libreoffice-help-fi

- libreoffice-help-fr

- libreoffice-help-gl

- libreoffice-help-gu

- libreoffice-help-he

- libreoffice-help-hi

- libreoffice-help-hr

- libreoffice-help-hu

- libreoffice-help-id

- libreoffice-help-it

- libreoffice-help-ja

- libreoffice-help-ko

- libreoffice-help-lt

- libreoffice-help-lv

- libreoffice-help-nb

- libreoffice-help-nl

- libreoffice-help-nn

- libreoffice-help-pl

- libreoffice-help-pt-BR

- libreoffice-help-pt-PT

- libreoffice-help-ro

- libreoffice-help-ru

- libreoffice-help-si

- libreoffice-help-sk

- libreoffice-help-sl

- libreoffice-help-sv

- libreoffice-help-ta

- libreoffice-help-tr

- libreoffice-help-uk

- libreoffice-help-zh-Hans

- libreoffice-help-zh-Hant

- libreoffice-impress

- libreoffice-langpack-af

- libreoffice-langpack-ar

- libreoffice-langpack-as

- libreoffice-langpack-bg

- libreoffice-langpack-bn

- libreoffice-langpack-br

- libreoffice-langpack-ca

- libreoffice-langpack-cs

- libreoffice-langpack-cy

- libreoffice-langpack-da

- libreoffice-langpack-de

- libreoffice-langpack-dz

- libreoffice-langpack-el

- libreoffice-langpack-en

- libreoffice-langpack-eo

- libreoffice-langpack-es

- libreoffice-langpack-et

- libreoffice-langpack-eu

- libreoffice-langpack-fa

- libreoffice-langpack-fi

- libreoffice-langpack-fr

- libreoffice-langpack-fy

- libreoffice-langpack-ga

- libreoffice-langpack-gl

- libreoffice-langpack-gu

- libreoffice-langpack-he

- libreoffice-langpack-hi

- libreoffice-langpack-hr

- libreoffice-langpack-hu

- libreoffice-langpack-id

- libreoffice-langpack-it

- libreoffice-langpack-ja

- libreoffice-langpack-kk

- libreoffice-langpack-kn

- libreoffice-langpack-ko

- libreoffice-langpack-lt
- libreoffice-langpack-lv
- libreoffice-langpack-mai
- libreoffice-langpack-ml
- libreoffice-langpack-mr
- libreoffice-langpack-nb
- libreoffice-langpack-nl
- libreoffice-langpack-nn
- libreoffice-langpack-nr
- libreoffice-langpack-nso
- libreoffice-langpack-or
- libreoffice-langpack-pa
- libreoffice-langpack-pl
- libreoffice-langpack-pt-BR
- libreoffice-langpack-pt-PT
- libreoffice-langpack-ro
- libreoffice-langpack-ru
- libreoffice-langpack-si
- libreoffice-langpack-sk
- libreoffice-langpack-sl
- libreoffice-langpack-sr
- libreoffice-langpack-ss
- libreoffice-langpack-st
- libreoffice-langpack-sv
- libreoffice-langpack-ta
- libreoffice-langpack-te
- libreoffice-langpack-th
- libreoffice-langpack-tn
- libreoffice-langpack-tr

- libreoffice-langpack-ts

- libreoffice-langpack-uk

- libreoffice-langpack-ve

- libreoffice-langpack-xh

- libreoffice-langpack-zh-Hans

- libreoffice-langpack-zh-Hant

- libreoffice-langpack-zu

- libreoffice-math

- libreoffice-ogltrans

- libreoffice-opensymbol-fonts

- libreoffice-pdfimport

- libreoffice-pyuno

- libreoffice-sdk

- libreoffice-sdk-doc

- libreoffice-ure

- libreoffice-ure-common

- libreoffice-wiki-publisher

- libreoffice-writer

- libreoffice-x11

- libreoffice-xsltfilter

- libreofficekit

- libsoup

- libsoup-devel

- libuser

- libuser-devel

- libwpe

- libwpe-devel

- mcpp

- mod_auth_mellon

- motif

- motif-devel

- pmdk-convert

- pmempool

- python3-pytz

- qla4xxx

- qt5

- qt5-assistant

- qt5-designer

- qt5-devel

- qt5-doctools

- qt5-linguist

- qt5-qdbusviewer

- qt5-qt3d

- qt5-qt3d-devel

- qt5-qt3d-doc

- qt5-qt3d-examples

- qt5-qtbase

- qt5-qtbase-common

- qt5-qtbase-devel

- qt5-qtbase-doc

- qt5-qtbase-examples

- qt5-qtbase-gui

- qt5-qtbase-mysql

- qt5-qtbase-odbc

- qt5-qtbase-postgresql

- qt5-qtbase-private-devel

- qt5-qtbase-static

- qt5-qtconnectivity

- qt5-qtconnectivity-devel

- qt5-qtconnectivity-doc

- qt5-qtconnectivity-examples

- qt5-qtdeclarative

- qt5-qtdeclarative-devel

- qt5-qtdeclarative-doc

- qt5-qtdeclarative-examples

- qt5-qtdeclarative-static

- qt5-qtdoc

- qt5-qtgraphicaleffects

- qt5-qtgraphicaleffects-doc

- qt5-qtimageformats

- qt5-qtimageformats-doc

- qt5-qtlocation

- qt5-qtlocation-devel

- qt5-qtlocation-doc

- qt5-qtlocation-examples

- qt5-qtmultimedia

- qt5-qtmultimedia-devel

- qt5-qtmultimedia-doc

- qt5-qtmultimedia-examples

- qt5-qtquickcontrols

- qt5-qtquickcontrols-doc

- qt5-qtquickcontrols-examples

- qt5-qtquickcontrols2

- qt5-qtquickcontrols2-devel

- qt5-qtquickcontrols2-doc

- qt5-qtquickcontrols2-examples

- qt5-qtscript

- qt5-qtscript-devel

- qt5-qtscript-doc

- qt5-qtscript-examples

- qt5-qtsensors

- qt5-qtsensors-devel

- qt5-qtsensors-doc

- qt5-qtsensors-examples

- qt5-qtserialbus

- qt5-qtserialbus-devel

- qt5-qtserialbus-doc

- qt5-qtserialbus-examples

- qt5-qtserialport

- qt5-qtserialport-devel

- qt5-qtserialport-doc

- qt5-qtserialport-examples

- qt5-qtsvg

- qt5-qtsvg-devel

- qt5-qtsvg-doc

- qt5-qtsvg-examples

- qt5-qttools

- qt5-qttools-common

- qt5-qttools-devel

- qt5-qttools-doc

- qt5-qttools-examples

- qt5-qttools-libs-designer

- qt5-qttools-libs-designercomponents

- qt5-qttools-libs-help

- qt5-qttools-static

- qt5-qttranslations

- qt5-qtwayland

- qt5-qtwayland-devel

- qt5-qtwayland-doc

- qt5-qtwayland-examples

- qt5-qtwebchannel

- qt5-qtwebchannel-devel

- qt5-qtwebchannel-doc

- qt5-qtwebchannel-examples

- qt5-qtwebsockets

- qt5-qtwebsockets-devel

- qt5-qtwebsockets-doc

- qt5-qtwebsockets-examples

- qt5-qtx11extras

- qt5-qtx11extras-devel

- qt5-qtx11extras-doc

- qt5-qtxmlpatterns

- qt5-qtxmlpatterns-devel

- qt5-qtxmlpatterns-doc

- qt5-qtxmlpatterns-examples

- qt5-rpm-macros

- qt5-srpm-macros

- team

- tigervnc

- tigervnc-icons

- tigervnc-license

- tigervnc-selinux

- tigervnc-server

- tigervnc-server-minimal

- tigervnc-server-module

- webkit2gtk3

- webkit2gtk3-devel

- webkit2gtk3-jsc

- webkit2gtk3-jsc-devel

- wpebackend-fdo

- wpebackend-fdo-devel

- xorg-x11-server-Xorg

- yp-tools

- ypbind

- ypserv

CHAPTER 11. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 9.4.

11.1. INSTALLER AND IMAGE CREATION

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

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

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

Bugzilla:1640697[1]

The **reboot --kexec** and **inst.kexec** commands do not provide a predictable system state

Performing a RHEL installation with the **reboot --kexec** Kickstart command or the **inst.kexec** kernel boot parameters do not provide the same predictable system state as a full reboot. As a consequence, switching to the installed system without rebooting can produce unpredictable results.

Note that the **kexec** feature is deprecated and will be removed in a future release of Red Hat Enterprise Linux.

Bugzilla:1697896[1]

Unexpected SELinux policies on systems where Anaconda is running as an application

When Anaconda is running as an application on an already installed system (for example to perform another installation to an image file using the **–image** anaconda option), the system is not prohibited to modify the SELinux types and attributes during installation. As a consequence, certain elements of SELinux policy might change on the system where Anaconda is running.

To work around this problem, do not run Anaconda on the production system. Instead, run Anaconda in a temporary virtual machine to keep the SELinux policy unchanged on a production system. Running anaconda as part of the system installation process such as installing from **boot.iso** or **dvd.iso** is not affected by this issue.

Bugzilla:2050140

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

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

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

As a workaround, use either of the following solution:

- When booting the installation, click the **Tab** key to edit the kernel command line, and change the boot option **inst.stage2=** to **inst.repo=**.

- To create a bootable USB device on Windows, use Fedora Media Writer.

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

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

Bugzilla:1877697[1]

The USB CD-ROM drive is not available as an installation source in Anaconda

Installation fails when the USB CD-ROM drive is the source for it and the Kickstart **ignoredisk --only-use=** command is specified. In this case, Anaconda cannot find and use this source disk.

To work around this problem, use the **harddrive --partition=sdX --dir=/** command to install from USB CD-ROM drive. As a result, the installation does not fail.

Jira:RHEL-4707

Hard drive partitioned installations with iso9660 filesystem fails

You cannot install RHEL on systems where the hard drive is partitioned with the **iso9660** filesystem. This is due to the updated installation code that is set to ignore any hard disk containing a **iso9660** file system partition. This happens even when RHEL is installed without using a DVD.

To workaround this problem, add the following script in the Kickstart file to format the disc before the installation starts.

Note: Before performing the workaround, backup the data available on the disk. The **wipefs** command formats all the existing data from the disk.

```
%pre
wipefs -a /dev/sda
%end
```

As a result, installations work as expected without any errors.

Jira:RHEL-4711

Anaconda fails to verify existence of an administrator user account

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

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

Bugzilla:2047713

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

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

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

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

Bugzilla:1997832[1]

RHEL for Edge installer image fails to create mount points when installing an rpm-ostree payload

When deploying **rpm-ostree** payloads, used for example in a RHEL for Edge installer image, the installer does not properly create some mount points for custom partitions. As a consequence, the installation is aborted with the following error:

> The command 'mount --bind /mnt/sysimage/data /mnt/sysroot/data' exited with the code 32.

To work around this issue:

- Use an automatic partitioning scheme and do not add any mount points manually.

- Manually assign mount points only inside **/var** directory. For example, /**var**/*my-mount-point*), and the following standard directories: /, **/boot**, **/var**.

As a result, the installation process finishes successfully.

Jira:RHEL-4741

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

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

As a workaround, do one of the following:

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

  a. List all connection profiles:

     ```
     # nmcli connection show
     ```

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

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

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

- Disable the auto connect network feature in Anaconda if no specific **ip=** or Kickstart network configuration is set.

  a. In the Anaconda GUI, navigate to **Network & Host Name**

  b. Select a network device to disable.

c. Click **Configure**.

d. On the **General** tab, clear the **Connect automatically with priority** checkbox.

e. Click **Save**.

Bugzilla:2115783[1]

Kickstart installations fail to configure the network connection

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

To work around this problem:

- Configure the network, for example using the **nmcli** tool, as a part of the **%pre** script.

- Use the installer boot options to configure the network for the **%pre** script.

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

Bugzilla:2173992

Enabling the FIPS mode is not supported when building rpm-ostree images with RHEL image builder

Currently, there is no support to enable the FIPS mode when building **rpm-ostree** images with RHEL image builder.

Jira:RHEL-4655

Images built with the stig profile remediation fails to boot with FIPS error

FIPS mode is not supported by RHEL image builder. When using RHEL image builder customized with the **xccdf_org.ssgproject.content_profile_stig** profile remediation, the system fails to boot with the following error:

> Warning: /boot//.vmlinuz-<kernel version>.x86_64.hmac does not exist
> FATAL: FIPS integrity test failed
> Refusing to continue

Enabling the FIPS policy manually after the system image installation with the **fips-mode-setup -- enable** command does not work, because the /**boot** directory is on a different partition. System boots successfully if FIPS is disabled. Currently, there is no workaround available.

> **NOTE**
>
> You can manually enable FIPS after installing the image by using the **fips-mode-setup -- enable** command.

Jira:RHEL-4649

Driver disk menu fails to display user inputs on the console

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

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

Jira:RHEL-4737

Kickstart installation fails due to missing packages with `systemd` service files in `%packages` section

If the Kickstart file uses the **services --enabled=…** directive to enable **systemd** services and packages containing the specified service file are not included in the **%packages** section, the RHEL installation process fails with the following error:

> Error enabling service <name_of_the_service>

To work around this problem, include the respective package with the service file in Kickstart's **%packages** section. As a result, RHEL installation completes, enabling expected services during installation.

Jira:RHEL-9633[1]

bootc-image-builder does not support building images from private registries

Currently, you cannot build base disk images which come from private registries by using **bootc-image-builder**. To workaround this issue, copy the private registry into your localhost, then build the image with the following arguments:

- **--local**

- **localhost/<image name>:tag** as the image

For example, to build your image:

```
sudo podman run \
--rm \
-it \
--privileged \
--pull=newer \
--security-opt label=type:unconfined_t \
-v ./config.toml:/config.toml \
-v ./output:/output \
-v /var/lib/containers/storage:/var/lib/containers/storage \
registry.redhat.io/rhel9/bootc-image-builder:latest
--type qcow2 \
--local \
quay.io/<namespace>/<image>:<tag>
```

Jira:RHEL-34054

11.2. SECURITY

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

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

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

Bugzilla:1681178[1]

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

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

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

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

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

Bugzilla:1685470[1]

With a specific syntax, **scp** empties files copied to themselves

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

scp localhost:/myfile localhost:/myfile

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

The problem has been fixed for the following syntaxes:

- **scp /myfile localhost:/myfile**

- **scp localhost:~/myfile ~/myfile**

Bugzilla:2056884

The OSCAP Anaconda add-on does not fetch tailored profiles in the graphical installation

The OSCAP Anaconda add-on does not provide an option to select or deselect tailoring of security profiles in the RHEL graphical installation. Starting from RHEL 8.8, the add-on does not take tailoring into account by default when installing from archives or RPM packages. Consequently, the installation displays the following error message instead of fetching an OSCAP tailored profile:

> There was an unexpected problem with the supplied content.

To work around this problem, you must specify paths in the **%addon org_fedora_oscap** section of your Kickstart file, for example:

> xccdf-path = /usr/share/xml/scap/sc_tailoring/ds-combined.xml
> tailoring-path = /usr/share/xml/scap/sc_tailoring/tailoring-xccdf.xml

As a result, you can use the graphical installation for OSCAP tailored profiles only with the corresponding Kickstart specifications.

Jira:RHEL-1824

Ansible remediations require additional collections

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

For an Ansible remediation, perform the following steps:

1. Install the required packages:

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

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

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

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

   > # ANSIBLE_COLLECTIONS_PATH=/usr/share/rhc-worker-playbook/ansible/collections/ansible_collections/ ansible-playbook -c local -i localhost, rhel9-playbook-*cis_server_l1*.yml

   Replace *cis_server_l1* with the ID of the profile against which you want to remediate the system.

As a result, the Ansible content is processed correctly.

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

Jira:RHEL-1800

Keylime does not accept concatenated PEM certificates

When Keylime receives a certificate chain as multiple certificates in the PEM format concatenated in a single file, the **keylime-agent-rust** Keylime component does not correctly use all the provided certificates during signature verification, resulting in a TLS handshake failure. As a consequence, the

client components (**keylime_verifier** and **keylime_tenant**) cannot connect to the Keylime agent. To work around this problem, use just one certificate instead of multiple certificates.

Jira:RHELPLAN-157225[1]

Keylime refuses runtime policies whose digests start with a backslash

The current script for generating runtime policies, **create_runtime_policy.sh**, uses SHA checksum functions, for example, **sha256sum**, to compute the file digest. However, when the input file name contains a backslash or **\n**, the checksum function adds a backslash before the digest in its output. In such cases, the generated policy file is malformed. When provided with the malformed policy file, the Keylime tenant produces the following or similar error message: **me.tenant - ERROR - Response code 400: Runtime policy is malformatted**. To work around the problem, remove the backslash from the malformed policy file manually by entering the following command: **sed -i 's/^\\//g' <malformed_file_name>**.

Jira:RHEL-11867[1]

Keylime agent rejects requests from the verifier after update

When the API version number of the Keylime agent (**keylime-agent-rust**) has been updated, the agent rejects requests that use a different version. As a consequence, if a Keylime agent is added to a verifier and then updated, the verifier tries to contact the agent using the old API version. The agent rejects this request and fails the attestation. To work around this problem, update the verifier (**keylime-verifier**) before updating the agent (**keylime-agent-rust**). As a result, when the agents are updated, the verifier detects the API change and updates its stored data accordingly.

Jira:RHEL-1518[1]

Missing files in **trustdb** cause denials for **fapolicyd**

When **fapolicyd** is installed with the Ansible DISA STIG profile, a race condition causes the **trustdb** database to be out of sync with the **rpmdb** database. As a consequence, missing files in **trustdb** cause denials on the system. To work around this problem, restart **fapolicyd** or run the Ansible DISA STIG profile again.

Jira:RHEL-24345[1]

The **fapolicyd** utility incorrectly allows executing changed files

Correctly, the IMA hash of a file should update after any change to the file, and **fapolicyd** should prevent execution of the changed file. However, this does not happen due to differences in IMA policy setup and in file hashing by the **evctml** utility. As a result, the IMA hash is not updated in the extended attribute of a changed file. Consequently, **fapolicyd** incorrectly allows the execution of the changed file.

Jira:RHEL-520[1]

Default SELinux policy allows unconfined executables to make their stack executable

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

Bugzilla:2064274

SSH timeout rules in STIG profiles configure incorrect options

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

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

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

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

> Title: Set SSH Client Alive Count Max to zero
> CCE Identifier: CCE-90271-8
> Rule ID: xccdf_org.ssgproject.content_rule_sshd_set_keepalive_0
>
> Title: Set SSH Idle Timeout Interval
> CCE Identifier: CCE-90811-1
> Rule ID: xccdf_org.ssgproject.content_rule_sshd_set_idle_timeout

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

Bugzilla:2038978

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

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

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

Bugzilla:2070722

OpenSCAP memory-consumption problems

On systems with limited memory, the OpenSCAP scanner might stop prematurely or it might not generate the results files. To work around this problem, you can customize the scanning profile to deselect rules that involve recursion over the entire / file system:

- **rpm_verify_hashes**

- **rpm_verify_permissions**

- **rpm_verify_ownership**

- **file_permissions_unauthorized_world_writable**

- **no_files_unowned_by_user**

- **dir_perms_world_writable_system_owned**

- **file_permissions_unauthorized_suid**

- **file_permissions_unauthorized_sgid**

- **file_permissions_ungroupowned**

- **dir_perms_world_writable_sticky_bits**

For more details and more workarounds, see the related Knowledgebase article.

Bugzilla:2161499

Remediating service-related rules during kickstart installations might fail

During a kickstart installation, the OpenSCAP utility sometimes incorrectly shows that a service **enable** or **disable** state remediation is not needed. Consequently, OpenSCAP might set the services on the installed system to a non-compliant state. As a workaround, you can scan and remediate the system after the kickstart installation. This will fix the service-related issues.

Jira:RHELPLAN-44202[1]

Interoperability of FIPS:OSPP hosts impacted due to CNSA 1.0

The **OSPP** subpolicy has been aligned with Commercial National Security Algorithm (CNSA) 1.0. This affects the interoperability of hosts that use the **FIPS:OSPP** policy-subpolicy combination, with the following major aspects:

- Minimum RSA key size is mandated at 3072 bits.

- Algorithm negotiations no longer support AES-128 ciphers, the secp256r1 elliptic curve, and the FFDHE-2048 group.

Jira:RHEL-2735[1]

Missing rules in the SELinux policy block permissions to SQL databases

Missing permission rules from the SELinux policy block connections to SQL databases. Consequently, the FIDO Device Onboard (FDO) services **fdo-manufacturing-server.service**, **fdo-owner-onboarding-server.service**, and **fdo-rendezvous-server.service** cannot connect with FDO databases, such as PostgreSQL and SQLite. Therefore, the system cannot start the FDO by using the supported databases for credentials and other parameters, such as storing ownership vouchers.

You can work around this problem by performing the following steps:

1. Create a new file named **local_fdo_update.cil** and enter the missing SELinux policy rules:

   ```
   (allow fdo_t etc_t (file (write)))
   (allow fdo_t fdo_conf_t (file (append create rename setattr unlink write )))
   (allow fdo_t fdo_var_lib_t (dir (add_name remove_name write )))
   (allow fdo_t fdo_var_lib_t (file (create setattr unlink write )))
   (allow fdo_t krb5_keytab_t (dir (search)))
   (allow fdo_t postgresql_port_t (tcp_socket (name_connect)))
   (allow fdo_t sssd_t (unix_stream_socket (connectto)))
   (allow fdo_t sssd_var_run_t (sock_file (write)))
   ```

2. Install the policy module package:

   ```
   # semodule -i local_fdo_update.cil
   ```

As a consequence, FDO can connect with the PostgreSQL database and also fix problems related to SQLite permissions over **/var/lib/fdo/**, where the SQLite database files are expected to be located.

Jira:RHEL-28814

11.3. RHEL FOR EDGE

The open-vm-tools package is not available in the edge-vsphere image

Currently, the **open-vm-tools** package is not installed by default in the **edge-vsphere** image. To workaround this issue, include the package in the blueprint customization. When using the **edge-vsphere** image type, add the **open-vm-tools** in the blueprint for the RHEL for Edge Container image or the RHEL for Edge Commit image.

Jira:RHELDOCS-16574[1]

11.4. SOFTWARE MANAGEMENT

The Installation process sometimes becomes unresponsive

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

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

To workaround this problem, restart the installation process.

Bugzilla:2073510

Running createrepo_c on local repositories generates duplicate repodata files

When you run the **createrepo_c** command on local repositories, it generates duplicate copies of **repodata** files, one of the copies is compressed and one is not. There is no workaround available, however, you can safely ignore the duplicate files. The **createrepo_c** command generates duplicate copies because of requirements and differences in other tools relying on repositories created by using **createrepo_c**.

Bugzilla:2056318

11.5. SHELLS AND COMMAND-LINE TOOLS

Renaming network interfaces using ifcfg files fails

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

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

Bugzilla:2018112[1]

The chkconfig package is not installed by default in RHEL 9

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

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

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

Bugzilla:2053598[1]

The initscripts package is not installed by default

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

Bugzilla:2082303

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

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

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

```
# dnf install libxkbcommon
```

Jira:RHEL-6105

The %vmeff metric from the sysstat package displays incorrect values

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

Jira:RHEL-12009

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

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

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

Jira:RHEL-6995[1]

The ReaR rescue image on UEFI systems with Secure Boot enabled fails to boot with the default settings

ReaR image creation by using the **rear mkrescue** or **rear mkbackup** command fails with the following message:

> grub2-mkstandalone may fail to make a bootable EFI image of GRUB2 (no /usr/*/grub*/x86_64-efi/moddep.lst file)
> (...)
> grub2-mkstandalone: error: /usr/lib/grub/x86_64-efi/modinfo.sh doesn't exist. Please specify --target or --directory.

The missing files are part of the **grub2-efi-x64-modules** package. If you install this package, the rescue image is created successfully without any errors. When the **UEFI** Secure Boot is enabled, the rescue image is not bootable because it uses a boot loader that is not signed.

To work around this problem, add the following variables to the **/etc/rear/local.conf** or **/etc/rear/site.conf** ReaR configuration file):

> UEFI_BOOTLOADER=/boot/efi/EFI/redhat/grubx64.efi
> SECURE_BOOT_BOOTLOADER=/boot/efi/EFI/redhat/shimx64.efi

With the suggested workaround, the image can be produced successfully even on systems without the **grub2-efi-x64-modules** package, and it is bootable on systems with Secure Boot enabled. In addition, during the system recovery, the bootloader of the recovered system is set to the **EFI** shim bootloader.

For more information about **UEFI**, **Secure Boot**, and **shim bootloader**, see the UEFI: what happens when booting the system Knowledge Base article.

Jira:RHELDOCS-18064[1]

The %util column produced by sar and iostat utilities is invalid

When you collect system usage statistics by using the **sar** or **iostat** utilities, the **%util** column produced by **sar** or **iostat** might contain invalid data.

Jira:RHEL-26275[1]

The lsb-release binary is not available in RHEL 9

The information in **/etc/os-release** was previously available by calling the **lsb-release** binary. This binary was included in the **redhat-lsb package**, which was removed in RHEL 9. Now, you can display information about the operating system, such as the distribution, version, code name, and associated metadata, by reading the **/etc/os-release** file. This file is provided by Red Hat and any changes to it will be overwritten with each update of the **redhat-release** package. The format of the file is **KEY=VALUE**, and you can safely source the data for a shell script.

Jira:RHELDOCS-16427[1]

11.6. INFRASTRUCTURE SERVICES

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

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

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

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

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

Bugzilla:2070495

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

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

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

Bugzilla:1984982

libotr is not compliant with FIPS

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

Bugzilla:2086562

11.7. NETWORKING

kTLS does not support offloading of TLS 1.3 to NICs

Kernel Transport Layer Security (kTLS) does not support offloading of TLS 1.3 to NICs. Consequently, software encryption is used with TLS 1.3 even when the NICs support TLS offload. To work around this problem, disable TLS 1.3 if offload is required. As a result, you can offload only TLS 1.2. When TLS 1.3 is in use, there is lower performance, since TLS 1.3 cannot be offloaded.

Bugzilla:2000616[1]

Failure to update the session key causes the connection to break

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

Bugzilla:2013650[1]

11.8. KERNEL

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

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

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

> <jemalloc>: Unsupported system page size

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

- Use the appropriate build configuration or environment options to create 4k and 64k page size compatible binaries.

- Build any user space packages that use **jemalloc** after booting into the final 64k kernel and runtime environment.

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

> # cargo install fd-find --force

Bugzilla:2167783[1]

Upgrading to the latest real-time kernel with **dnf** does not install multiple kernel versions in parallel

Installing the latest real-time kernel with the **dnf** package manager requires resolving package dependencies to retain the new and current kernel versions simultaneously. By default, **dnf** removes the older **kernel-rt** package during the upgrade.

As a workaround, add the current **kernel-rt** package to the **installonlypkgs** option in the **/etc/yum.conf** configuration file, for example, **installonlypkgs=kernel-rt**.

The **installonlypkgs** option appends **kernel-rt** to the default list used by **dnf**. Packages listed in **installonlypkgs** directive are not removed automatically and therefore support multiple kernel versions to install simultaneously.

Note that having multiple kernels installed is a way to have a fallback option when working with a new kernel version.

Bugzilla:2181571[1]

The **Delay Accounting** functionality does not display the **SWAPIN** and **IO%** statistics columns by default

The **Delayed Accounting** functionality, unlike early versions, is disabled by default. Consequently, the **iotop** application does not show the **SWAPIN** and **IO%** statistics columns and displays the following warning:

> CONFIG_TASK_DELAY_ACCT not enabled in kernel, cannot determine SWAPIN and IO%

The **Delay Accounting** functionality, using the **taskstats** interface, provides the delay statistics for all tasks or threads that belong to a thread group. Delays in task execution occur when they wait for a kernel resource to become available, for example, a task waiting for a free CPU to run on. The statistics help in setting a task's CPU priority, I/O priority, and **rss** limit values appropriately.

As a workaround, you can enable the **delayacct** boot option either at run time or boot.

- To enable **delayacct** at run time, enter:

  > echo 1 > /proc/sys/kernel/task_delayacct

  Note that this command enables the feature system wide, but only for the tasks that you start after running this command.

- To enable **delayacct** permanently at boot, use one of the following procedures:

  - Edit the **/etc/sysctl.conf** file to override the default parameters:

    a. Add the following entry to the **/etc/sysctl.conf** file:

       > kernel.task_delayacct = 1

       For more information, see How to set sysctl variables on Red Hat Enterprise Linux .

    b. Reboot the system for changes to take effect.

  - Add the **delayacct** option to the kernel command line.
    For more information, see Configuring kernel command-line parameters.

As a result, the **iotop** application displays the **SWAPIN** and **IO%** statistics columns.

Bugzilla:2132480[1]

Hardware certification of the real-time kernel on systems with large core-counts might require passing the **skew-tick=1** boot parameter

Large or moderate sized systems with numerous sockets and large core-counts can experience latency spikes due to lock contentions on **xtime_lock**, which is used in the timekeeping system. As a consequence, latency spikes and delays in hardware certifications might occur on multiprocessing systems. As a workaround, you can offset the timer tick per CPU to start at a different time by adding the **skew_tick=1** boot parameter.

To avoid lock conflicts, enable **skew_tick=1**:

1. Enable the **skew_tick=1** parameter with **grubby**.

   > # grubby --update-kernel=ALL --args="skew_tick=1"

2. Reboot for changes to take effect.

3. Verify the new settings by displaying the kernel parameters you pass during boot.

```
cat /proc/cmdline
```

Note that enabling **skew_tick=1** causes a significant increase in power consumption and, therefore, it must be enabled only if you are running latency sensitive real-time workloads.

Jira:RHEL-9318[1]

The **kdump** mechanism fails to capture the **vmcore** file on LUKS-encrypted targets

When running **kdump** on systems with Linux Unified Key Setup (LUKS) encrypted partitions, systems require a certain amount of available memory. When the available memory is less than the required amount of memory, the **systemd-cryptsetup** service fails to mount the partition. Consequently, the second kernel fails to capture the crash dump file on the LUKS-encrypted targets.

As a workaround, query the **Recommended crashkernel value** and gradually increase the memory size to an appropriate value. The **Recommended crashkernel value** can serve as reference to set the required memory size.

1. Print the estimate crash kernel value.

```
# kdumpctl estimate
```

2. Configure the amount of required memory by increasing the **crashkernel** value.

```
# grubby --args=crashkernel=652M --update-kernel=ALL
```

3. Reboot the system for changes to take effect.

```
# reboot
```

As a result, **kdump** works correctly on systems with LUKS-encrypted partitions.

Jira:RHEL-11196[1]

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

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

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

As a workaround, use one of the following solutions:

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

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

- Copy the **znet** configuration information from the inactive connection profile to the active connection profile:

a. Run the **nmcli** command to query the **NetworkManager** connection profiles:

```
# nmcli connection show

NAME                  UUID            TYPE   Device

bridge-br0          ed391a43-bdea-4170-b8a2 bridge   br0
bridge-slave-enc600  caf7f770-1e55-4126-a2f4 ethernet enc600
enc600              bc293b8d-ef1e-45f6-bad1 ethernet --
```

b. Update the active profile with configuration information from the inactive connection:

```
#!/bin/bash
 inactive_connection=enc600
 active_connection=bridge-slave-enc600
 for name in nettype subchannels options; do
 field=802-3-ethernet.s390-$name
 val=$(nmcli --get-values "$field"connection show "$inactive_connection")
 nmcli connection modify "$active_connection" "$field" $val"
 done
```

c. Restart the **kdump** service for changes to take effect:

```
# kdumpctl restart
```

Bugzilla:2064708

The **iwl7260-firmware** breaks Wi-Fi on Intel Wi-Fi 6 AX200, AX210, and Lenovo ThinkPad P1 Gen 4

After updating the **iwl7260-firmware** or **iwl7260-wifi** driver to the version provided by RHEL 9.1 and later, the hardware gets into an incorrect internal state. reports its state incorrectly. Consequently, Intel Wifi 6 cards may not work and display the error message:

```
kernel: iwlwifi 0000:09:00.0: Failed to start RT ucode: -110
kernel: iwlwifi 0000:09:00.0: WRT: Collecting data: ini trigger 13 fired (delay=0ms)
kernel: iwlwifi 0000:09:00.0: Failed to run INIT ucode: -110
```

An unconfirmed workaround is to power off the system and back on again. Do not reboot.

Bugzilla:2129288[1]

weak-modules from **kmod** fails to work with module inter-dependencies

The **weak-modules** script provided by the **kmod** package determines which modules are kABI-compatible with installed kernels. However, while checking modules' kernel compatibility, **weak-modules** processes modules symbol dependencies from higher to lower release of the kernel for which they were built. As a consequence, modules with inter-dependencies built against different kernel releases might be interpreted as non-compatible, and therefore the **weak-modules** script fails to work in this scenario.

To work around the problem, build or put the extra modules against the latest stock kernel before you install the new kernel.

Bugzilla:2103605[1]

The Intel® **i40e** adapter permanently fails on IBM Power10

When the **i40e** adapter encounters an I/O error on IBM Power10 systems, the Enhanced I/O Error Handling (EEH) kernel services trigger the network driver's reset and recovery. However, EEH repeatedly reports I/O errors until the **i40e** driver reaches the predefined maximum of EEH freezes. As a consequence, EEH causes the device to fail permanently.

Jira:RHEL-15404[1]

dkms provides an incorrect warning on program failure with correctly compiled drivers on 64-bit ARM CPUs

The Dynamic Kernel Module Support (**dkms**) utility does not recognize that the kernel headers for 64-bit ARM CPUs work for both the kernels with 4 kilobytes and 64 kilobytes page sizes. As a result, when the kernel update is performed and the **kernel-64k-devel** package is not installed, **dkms** provides an incorrect warning on why the program failed on correctly compiled drivers. To work around this problem, install the **kernel-headers** package, which contains header files for both types of ARM CPU architectures and is not specific to **dkms** and its requirements.

Jira:RHEL-25967[1]

11.9. FILE SYSTEMS AND STORAGE

Device Mapper Multipath is not supported with NVMe/TCP

Using Device Mapper Multipath with the **nvme-tcp** driver can result in the Call Trace warnings and system instability. To work around this problem, NVMe/TCP users must enable native NVMe multipathing and not use the **device-mapper-multipath** tools with NVMe.

By default, Native NVMe multipathing is enabled in RHEL 9. For more information, see Enabling multipathing on NVMe devices.

Bugzilla:2033080[1]

The **blk-availability** systemd service deactivates complex device stacks

In **systemd**, the default block deactivation code does not always handle complex stacks of virtual block devices correctly. In some configurations, virtual devices might not be removed during the shutdown, which causes error messages to be logged. To work around this problem, deactivate complex block device stacks by executing the following command:

```
# systemctl enable --now blk-availability.service
```

As a result, complex virtual device stacks are correctly deactivated during shutdown and do not produce error messages.

Bugzilla:2011699[1]

Disabling quota accounting is no longer possible for an XFS filesystem mounted with quotas enabled

Starting with RHEL 9.2, it is no longer possible to disable quota accounting on an XFS filesystem which has been mounted with quotas enabled.

To work around this issue, disable quota accounting by remounting the filesystem, with the quota option removed.

Bugzilla:2160619[1]

udev rule change for NVMe devices

There is a udev rule change for NVMe devices that adds **OPTIONS="string_escape=replace"** parameter. This leads to a disk by-id naming change for some vendors, if the serial number of your device has leading whitespace.

Bugzilla:2185048

NVMe/FC devices cannot be reliably used in a Kickstart file

NVMe/FC devices can be unavailable during parsing or execution of pre-scripts of the Kickstart file, which can cause the Kickstart installation to fail. To work around this issue, update the boot argument to **inst.wait_for_disks=30**. This option causes a delay of 30 seconds, and should provide enough time for the NVMe/FC device to connect. With this workaround along with the NVMe/FC devices connecting in time, the Kickstart installation proceeds without issues.

Jira:RHEL-8164[1]

Kernel panic while using the qedi driver

While using the **qedi** iSCSI driver, the kernel panics after OS boots. To work around this issue, disable the **kfence** runtime memory error detector feature by adding **kfence.sample_interval=0** to the kernel boot command line.

Jira:RHEL-8466[1]

ARM-based systems fail to update with a 64k page size kernel when vdo is installed

While installing the **vdo** package, RHEL installs the **kmod-kvdo** package and a kernel with **4k** page size as dependencies. As a consequence, updates from RHEL 9.3 to 9.x fail because **kmod-kvdo** conflicts with the 64k kernel. To work around this issue, remove the **vdo** package and its dependencies prior to attempting to update.

Jira:RHEL-8354

lldpad is auto enabled even for qedf adapters

When using a QLogic Corp. FastLinQ QL45000 Series 10/25/40/50GbE, FCOE Controller automatically enables the **lldpad** daemon on systems running RHV. As a consequence, I/O operations are aborted with an error, for example, **[qedf_eh_abort:xxxx]:1: Aborting io_req=ff5d85a9dcf3xxxx**.

To work around this problem, disableLink Layer Discovery Protocol (LLDP) and then enable it for interfaces that can be set on the **vdsm** configuration level. For more information, https://access.redhat.com/solutions/6963195.

Jira:RHEL-8104[1]

System fails to boot when iommu is enabled

By enabling the Input-Output Memory Management Unit (IOMMU) on AMD platforms when the BNX2I adapter is in use, a system fails to boot with the Direct Memory Access Remapping (DMAR) timeout errors. To work around this problem, disable the IOMMU before booting by using the kernel command-line option, **iommu=off**. As a result, the system boots without any errors.

Jira:RHEL-25730[1]

11.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

python3.11-lxml does not provide the lxml.isoschematron submodule

The **python3.11-lxml** package is distributed without the **lxml.isoschematron** submodule because it is not under an open source license. The submodule implements ISO Schematron support. As an alternative, pre-ISO-Schematron validation is available in the **lxml.etree.Schematron** class. The remaining content of the **python3.11-lxml** package is unaffected.

Bugzilla:2157708

The --ssl-fips-mode option in MySQL and MariaDB does not change FIPS mode

The **--ssl-fips-mode** option in **MySQL** and **MariaDB** in RHEL works differently than in upstream.

In RHEL 9, if you use **--ssl-fips-mode** as an argument for the **mysqld** or **mariadbd** daemon, or if you use **ssl-fips-mode** in the **MySQL** or **MariaDB** server configuration files, **--ssl-fips-mode** does not change FIPS mode for these database servers.

Instead:

- If you set **--ssl-fips-mode** to **ON**, the **mysqld** or **mariadbd** server daemon does not start.

- If you set **--ssl-fips-mode** to **OFF** on a FIPS-enabled system, the **mysqld** or **mariadbd** server daemons still run in FIPS mode.

This is expected because FIPS mode should be enabled or disabled for the whole RHEL system, not for specific components.

Therefore, do not use the **--ssl-fips-mode** option in **MySQL** or **MariaDB** in RHEL. Instead, ensure FIPS mode is enabled on the whole RHEL system:

- Preferably, install RHEL with FIPS mode enabled. Enabling FIPS mode during the installation ensures that the system generates all keys with FIPS-approved algorithms and continuous monitoring tests in place. For information about installing RHEL in FIPS mode, see Installing the system in FIPS mode.

- Alternatively, you can switch FIPS mode for the entire RHEL system by following the procedure in Switching the system to FIPS mode .

Bugzilla:1991500

11.11. IDENTITY MANAGEMENT

MIT Kerberos does not support ECC certificates for PKINIT

MIT Kerberos does not implement the RFC5349 request for comments document, which describes the design of elliptic-curve cryptography (ECC) support in Public Key Cryptography for initial authentication (PKINIT). Consequently, the MIT **krb5-pkinit** package, used by RHEL, does not support ECC certificates. For more information, see Elliptic Curve Cryptography (ECC) Support for Public Key Cryptography for Initial Authentication in Kerberos (PKINIT).

Jira:RHEL-4902

The DEFAULT:SHA1 subpolicy has to be set on RHEL 9 clients for PKINIT to work against AD KDCs

The SHA-1 digest algorithm has been deprecated in RHEL 9, and CMS messages for Public Key Cryptography for initial authentication (PKINIT) are now signed with the stronger SHA-256 algorithm.

However, the Active Directory (AD) Kerberos Distribution Center (KDC) still uses the SHA-1 digest algorithm to sign CMS messages. As a result, RHEL 9 Kerberos clients fail to authenticate users by using PKINIT against an AD KDC.

To work around the problem, enable support for the SHA-1 algorithm on your RHEL 9 systems with the following command:

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

Bugzilla:2060798

The PKINIT authentication of a user fails if a RHEL 9 Kerberos agent communicates with a non-RHEL-9 and non-AD Kerberos agent

If a RHEL 9 Kerberos agent, either a client or Kerberos Distribution Center (KDC), interacts with a non-RHEL-9 Kerberos agent that is not an Active Directory (AD) agent, the PKINIT authentication of the user fails. To work around the problem, perform one of the following actions:

- Set the RHEL 9 agent's crypto-policy to **DEFAULT:SHA1** to allow the verification of SHA-1 signatures:

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

- Update the non-RHEL-9 and non-AD agent to ensure it does not sign CMS data using the SHA-1 algorithm. For this, update your Kerberos client or KDC packages to the versions that use SHA-256 instead of SHA-1:

  - CentOS 9 Stream: krb5-1.19.1-15

  - RHEL 8.7: krb5-1.18.2-17

  - RHEL 7.9: krb5-1.15.1-53

  - Fedora Rawhide/36: krb5-1.19.2-7

  - Fedora 35/34: krb5-1.19.2-3

As a result, the PKINIT authentication of the user works correctly.

Note that for other operating systems, it is the krb5-1.20 release that ensures that the agent signs CMS data with SHA-256 instead of SHA-1.

See also The DEFAULT:SHA1 subpolicy has to be set on RHEL 9 clients for PKINIT to work against AD KDCs.

Jira:RHEL-4875

FIPS support for AD trust requires the AD-SUPPORT crypto subpolicy

Active Directory (AD) uses AES SHA-1 HMAC encryption types, which are not allowed in FIPS mode on RHEL 9 by default. If you want to use RHEL 9 IdM hosts with an AD trust, enable support for AES SHA-1 HMAC encryption types before installing IdM software.

Since FIPS compliance is a process that involves both technical and organizational agreements, consult your FIPS auditor before enabling the **AD-SUPPORT** subpolicy to allow technical measures to support AES SHA-1 HMAC encryption types, and then install RHEL IdM:

```
# update-crypto-policies --set FIPS:AD-SUPPORT
```

Bugzilla:2057471

Heimdal client fails to authenticate a user using PKINIT against RHEL 9 KDC

By default, a Heimdal Kerberos client initiates the PKINIT authentication of an IdM user by using Modular Exponential (MODP) Diffie-Hellman Group 2 for Internet Key Exchange (IKE). However, the MIT Kerberos Distribution Center (KDC) on RHEL 9 only supports MODP Group 14 and 16.

Consequently, the pre-autentication request fails with the **krb5_get_init_creds: PREAUTH_FAILED** error on the Heimdal client and **Key parameters not accepted** on the RHEL MIT KDC.

To work around this problem, ensure that the Heimdal client uses MODP Group 14. Set the **pkinit_dh_min_bits** parameter in the **libdefaults** section of the client configuration file to 1759:

```
[libdefaults]
pkinit_dh_min_bits = 1759
```

As a result, the Heimdal client completes the PKINIT pre-authentication against the RHEL MIT KDC.

Jira:RHEL-4889

IdM in FIPS mode does not support using the NTLMSSP protocol to establish a two-way cross-forest trust

Establishing a two-way cross-forest trust between Active Directory (AD) and Identity Management (IdM) with FIPS mode enabled fails because the New Technology LAN Manager Security Support Provider (NTLMSSP) authentication is not FIPS-compliant. IdM in FIPS mode does not accept the RC4 NTLM hash that the AD domain controller uses when attempting to authenticate.

Jira:RHEL-12154[1]

Users without SIDs cannot log in to IdM after an upgrade

After upgrading your IdM replica to RHEL 9.2, the IdM Kerberos Distribution Center (KDC) might fail to issue ticket-granting tickets (TGTs) to users who do not have Security Identifiers (SIDs) assigned to their accounts. Consequently, the users cannot log in to their accounts.

To work around the problem, generate SIDs by running the following command as an IdM administrator on another IdM replica in the topology:

```
# ipa config-mod --enable-sid --add-sids
```

Afterward, if users still cannot log in, examine the Directory Server error log. You might have to adjust ID ranges to include user POSIX identities.

See the When upgrading to RHEL9, IDM users are not able to login anymore   Knowledgebase solution for more information.

Jira:RHELPLAN-157939[1]

Migrated IdM users might be unable to log in due to mismatching domain SIDs

If you have used the **ipa migrate-ds** script to migrate users from one IdM deployment to another, those users might have problems using IdM services because their previously existing Security Identifiers (SIDs) do not have the domain SID of the current IdM environment. For example, those users can retrieve a Kerberos ticket with the **kinit** utility, but they cannot log in. To work around this problem, see the following Knowledgebase article: Migrated IdM users unable to log in due to mismatching domain SIDs.

Jira:RHELPLAN-109613[1]

MIT **krb5** user fails to obtain an AD TGT because of incompatible encryption types generating the user PAC

In MIT **krb5 1.20** and later packages, a Privilege Attribute Certificate (PAC) is included in all Kerberos tickets by default. The MIT Kerberos Distribution Center (KDC) selects the strongest encryption type available to generate the KDC checksum in the PAC, which currently is the **AES HMAC-SHA2** encryption types defined in RFC8009. However, Active Directory (AD) does not support this RFC. Consequently, in an AD-MIT cross-realm setup, an MIT **krb5** user fails to obtain an AD ticket-granting ticket (TGT) because the cross-realm TGT generated by MIT KDC contains an incompatible KDC checksum type in the PAC.

To work around the problem, set the **disable_pac** parameter to **true** for the MIT realm in the **[realms]** section of the **/var/kerberos/krb5kdc/kdc.conf** configuration file. As a result, the MIT KDC generates tickets without PAC, which means that AD skips the failing checksum verification and an MIT **krb5** user can obtain an AD TGT.

Bugzilla:2016312

Potential risk when using the default value for **ldap_id_use_start_tls** option

When using **ldap://** without TLS for identity lookups, it can pose a risk for an attack vector. Particularly a man-in-the-middle (MITM) attack which could allow an attacker to impersonate a user by altering, for example, the UID or GID of an object returned in an LDAP search.

Currently, the SSSD configuration option to enforce TLS, **ldap_id_use_start_tls**, defaults to **false**. Ensure that your setup operates in a trusted environment and decide if it is safe to use unencrypted communication for **id_provider = ldap**. Note **id_provider = ad** and **id_provider = ipa** are not affected as they use encrypted connections protected by SASL and GSSAPI.

If it is not safe to use unencrypted communication, enforce TLS by setting the **ldap_id_use_start_tls** option to **true** in the **/etc/sssd/sssd.conf** file. The default behavior is planned to be changed in a future release of RHEL.

Jira:RHELPLAN-155168[1]

Adding a RHEL 9 replica in FIPS mode to an IdM deployment in FIPS mode that was initialized with RHEL 8.6 or earlier fails

The default RHEL 9 FIPS cryptographic policy aiming to comply with FIPS 140-3 does not allow the use of the AES HMAC-SHA1 encryption types' key derivation function as defined by RFC3961, section 5.1.

This constraint is a blocker when adding a RHEL 9 Identity Management (IdM) replica in FIPS mode to a

RHEL 8 IdM environment in FIPS mode in which the first server was installed on a RHEL 8.6 system or earlier. This is because there are no common encryption types between RHEL 9 and the previous RHEL versions, which commonly use the AES HMAC-SHA1 encryption types but do not use the AES HMAC-SHA2 encryption types.

You can view the encryption type of your IdM master key by entering the following command on the server:

```
# kadmin.local getprinc K/M | grep -E '^Key:'
```

To work around the problem, enable the use of AES HMAC-SHA1 on the RHEL 9 replica:

```
update-crypto-policies --set FIPS:AD-SUPPORT
```

WARNING

This workaround might violate FIPS compliance.

As a result, adding the RHEL 9 replica to the IdM deployment proceeds correctly.

Note that there is ongoing work to provide a procedure to generate missing AES HMAC-SHA2-encrypted Kerberos keys on RHEL 7 and RHEL 8 servers. This will achieve FIPS 140-3 compliance on the RHEL 9 replica. However, this process will not be fully automated, because the design of Kerberos key cryptography makes it impossible to convert existing keys to different encryption types. The only way is to ask users to renew their passwords.

Jira:RHEL-4888

SSSD registers the DNS names properly

Previously, if the DNS was set up incorrectly, SSSD always failed the first attempt to register the DNS name. To work around the problem, this update provides a new parameter **dns_resolver_use_search_list**. Set **dns_resolver_use_search_list = false** to avoid using the DNS search list.

Bugzilla:1608496[1]

Installing a RHEL 7 IdM client with a RHEL 9.2 and later IdM server in FIPS mode fails due to EMS enforcement

The TLS **Extended Master Secret** (EMS) extension (RFC 7627) is now mandatory for TLS 1.2 connections on FIPS-enabled RHEL 9.2 and later systems. This is in accordance with FIPS-140-3 requirements. However, the **openssl** version available in RHEL 7.9 and lower does not support EMS. In consequence, installing a RHEL 7 Identity Management (IdM) client with a FIPS-enabled IdM server running on RHEL 9.2 and later fails.

If upgrading the host to RHEL 8 before installing an IdM client on it is not an option, work around the problem by removing the requirement for EMS usage on the RHEL 9 server by applying a NO-ENFORCE-EMS subpolicy on top of the FIPS crypto policy:

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

Note that this removal goes against the FIPS 140-3 requirements. As a result, you can establish and accept TLS 1.2 connections that do not use EMS, and the installation of a RHEL 7 IdM client succeeds.

Jira:RHEL-4955

The online backup and the online automembership rebuild tasks can acquire two locks resulting in a deadlock

If the online backup and the online automembership rebuild tasks attempt to acquire the same two locks in the opposite order, it can lead to an unrecoverable deadlock that requires you to stop and restart the server. To workaround this problem, do not launch the online backup and the online automembership rebuild tasks in parallel.

Jira:RHELDOCS-18065[1]

11.12. DESKTOP

VNC is not running after upgrading to RHEL 9

After upgrading from RHEL 8 to RHEL 9, the VNC server fails to start, even if it was previously enabled.

To work around the problem, manually enable the **vncserver** service after the system upgrade:

```
# systemctl enable --now vncserver@:port-number
```

As a result, VNC is now enabled and starts after every system boot as expected.

Bugzilla:2060308

User Creation screen is unresponsive

When installing RHEL using a graphical user interface, the User Creation screen is unresponsive. As a consequence, creating users during installation is more difficult.

To work around this problem, use one of the following solutions to create users:

- Run the installation in VNC mode and resize the VNC window.

- Create users after completing the installation process.

Jira:RHEL-11924[1]

WebKitGTK fails to display web pages on IBM Z

The WebKitGTK web browser engine fails when trying to display web pages on the IBM Z architecture. The web page remains blank and the WebKitGTK process ends unexpectedly.

As a consequence, you cannot use certain features of applications that use WebKitGTK to display web pages, such as the following:

- The Evolution mail client

- The GNOME Online Accounts settings

- The GNOME Help application

Jira:RHEL-4157

11.13. GRAPHICS INFRASTRUCTURES

NVIDIA drivers might revert to X.org

Under certain conditions, the proprietary NVIDIA drivers disable the Wayland display protocol and revert to the X.org display server:

- If the version of the NVIDIA driver is lower than 470.

- If the system is a laptop that uses hybrid graphics.

- If you have not enabled the required NVIDIA driver options.

Additionally, Wayland is enabled but the desktop session uses X.org by default if the version of the NVIDIA driver is lower than 510.

Jira:RHELPLAN-119001[1]

Night Light is not available on Wayland with NVIDIA

When the proprietary NVIDIA drivers are enabled on your system, the **Night Light** feature of GNOME is not available in Wayland sessions. The NVIDIA drivers do not currently support **Night Light**.

Jira:RHELPLAN-119852[1]

X.org configuration utilities do not work under Wayland

X.org utilities for manipulating the screen do not work in the Wayland session. Notably, the **xrandr** utility does not work under Wayland due to its different approach to handling, resolutions, rotations, and layout.

Jira:RHELPLAN-121049[1]

11.14. THE WEB CONSOLE

VNC console in the RHEL web console does not work correctly on ARM64

Currently, when you import a virtual machine (VM) in the RHEL web console on ARM64 architecture and then you try to interact with it in the VNC console, the console does not react to your input.

Additionally, when you create a VM in the web console on ARM64 architecture, the VNC console does not display the last lines of your input.

Jira:RHEL-31993[1]

11.15. RED HAT ENTERPRISE LINUX SYSTEM ROLES

If firewalld.service is masked, using the firewall RHEL System Role fails

If **firewalld.service** is masked on a RHEL system, the **firewall** RHEL System Role fails. To work around this problem, unmask the **firewalld.service**:

```
systemctl unmask firewalld.service
```

Bugzilla:2123859

Unable to register systems with environment names

The **rhc** system role fails to register the system when specifying environment names in **rhc_environment**. As a workaround, use environment IDs instead of environment names while registering.

Jira:RHEL-1172

Running Microsoft SQL Server 2022 in high-availability mode as an SELinux-confined application does not work

Microsoft SQL Server 2022 on RHEL 9.4 and later supports running as an SELinux-confined application. However, due to a limitation in Microsoft SQL Server, running the service as an SELinux-confined application does not work in high-availability mode. To work around this problem, you can run Microsoft SQL Server as an unconfined application if you require the service to be high available.

Note that this limitation also impacts installing Microsoft SQL Server when you use the **mssql** RHEL System Role to install this service.

Jira:RHELDOCS-17719[1]

The **mssql** RHEL System Role cannot configure Microsoft SQL Server with AD integration

The Microsoft SQL Server service does not provide the **adutil** tool that the service requires for the integration with Active Directory (AD). Consequently, you cannot use the **mssql** RHEL System Role to configure this scenario on a RHEL 9 managed node. No workaround is available, and you can use the RHEL System Role only to configure Microsoft SQL Server without AD integration on RHEL 9.

Jira:RHELDOCS-17720[1]

11.16. VIRTUALIZATION

Installing a virtual machine over https or ssh in some cases fails

Currently, the **virt-install** utility fails when attempting to install a guest operating system (OS) from an ISO source over a https or ssh connection - for example using **virt-install --cdrom https://example/path/to/image.iso**. Instead of creating a virtual machine (VM), the described operation ends unexpectedly with an **internal error: process exited while connecting to monitor** message.

Similarly, using the RHEL 9 web console to install a guest operating system fails and displays an **Unknown driver 'https'** error if you use an https or ssh URL, or the **Download OS** function.

To work around this problem, install **qemu-kvm-block-curl** and **qemu-kvm-block-ssh** on the host to enable https and ssh protocol support. Alternatively, use a different connection protocol or a different installation source.

Bugzilla:2014229

Using NVIDIA drivers in virtual machines disables Wayland

Currently, NVIDIA drivers are not compatible with the Wayland graphical session. As a consequence, RHEL guest operating systems that use NVIDIA drivers automatically disable Wayland and load an Xorg session instead. This primarily occurs in the following scenarios:

- When you pass through an NVIDIA GPU device to a RHEL virtual machine (VM)

- When you assign an NVIDIA vGPU mediated device to a RHEL VM

Jira:RHELPLAN-117234[1]

The **Milan** VM CPU type is sometimes not available on AMD Milan systems

On certain AMD Milan systems, the Enhanced REP MOVSB (**erms**) and Fast Short REP MOVSB (**fsrm**) feature flags are disabled in the BIOS by default. Consequently, the **Milan** CPU type might not be available on these systems. In addition, VM live migration between Milan hosts with different feature flag settings might fail. To work around these problems, manually turn on **erms** and **fsrm** in the BIOS of your host.

Bugzilla:2077767[1]

A **hostdev** interface with failover settings cannot be hot-plugged after being hot-unplugged

After removing a **hostdev** network interface with failover configuration from a running virtual machine (VM), the interface currently cannot be re-attached to the same running VM.

Jira:RHEL-7337

Live post-copy migration of VMs with failover VFs fails

Currently, attempting to post-copy migrate a running virtual machine (VM) fails if the VM uses a device with the virtual function (VF) failover capability enabled. To work around the problem, use the standard migration type, rather than post-copy migration.

Jira:RHEL-7335

Host network cannot ping VMs with VFs during live migration

When live migrating a virtual machine (VM) with a configured virtual function (VF), such as a VMs that uses virtual SR-IOV software, the network of the VM is not visible to other devices and the VM cannot be reached by commands such as **ping**. After the migration is finished, however, the problem no longer occurs.

Jira:RHEL-7336

Disabling AVX causes VMs to become unbootable

On a host machine that uses a CPU with Advanced Vector Extensions (AVX) support, attempting to boot a VM with AVX explicitly disabled currently fails, and instead triggers a kernel panic in the VM.

Bugzilla:2005173[1]

Windows VM fails to get IP address after network interface reset

Sometimes, Windows virtual machines fail to get an IP address after an automatic network interface reset. As a consequence, the VM fails to connect to the network. To work around this problem, disable and re-enable the network adapter driver in the Windows Device Manager.

Jira:RHEL-11366

Windows Server 2016 VMs sometimes stops working after hot-plugging a vCPU

Currently, assigning a vCPU to a running virtual machine (VM) with a Windows Server 2016 guest operating system might cause a variety of problems, such as the VM terminating unexpectedly, becoming unresponsive, or rebooting.

Bugzilla:1915715

Redundant error messages on VMs with NVIDIA passthrough devices

When using an Intel host machine with a RHEL 9.2 and later operating system, virtual machines (VMs) with a passed through NVDIA GPU device frequently log the following error message:

> Spurious APIC interrupt (vector 0xFF) on CPU#2, should never happen.

However, this error message does not impact the functionality of the VM and can be ignored. For details, see the Red Hat KnoweldgeBase .

Bugzilla:2149989[1]

Restarting the OVS service on a host might block network connectivity on its running VMs

When the Open vSwitch (OVS) service restarts or crashes on a host, virtual machines (VMs) that are running on this host cannot recover the state of the networking device. As a consequence, VMs might be completely unable to receive packets.

This problem only affects systems that use the packed virtqueue format in their **virtio** networking stack.

To work around this problem, use the **packed=off** parameter in the **virtio** networking device definition to disable packed virtqueue. With packed virtqueue disabled, the state of the networking device can, in some situations, be recovered from RAM.

Jira:RHEL-333

Recovering an interrupted post-copy VM migration might fail

If a post-copy migration of a virtual machine (VM) is interrupted and then immediately resumed on the same incoming port, the migration might fail with the following error: **Address already in use**

To work around this problem, wait at least 10 seconds before resuming the post-copy migration or switch to another port for migration recovery.

Jira:RHEL-7096

NUMA node mapping not working correctly on AMD EPYC CPUs

QEMU does not handle NUMA node mapping on AMD EPYC CPUs correctly. As a result, the performance of virtual machines (VMs) with these CPUs might be negatively impacted if using a NUMA node configuration. In addition, the VMs display a warning similar to the following during boot.

> sched: CPU #4's llc-sibling CPU #3 is not on the same node! [node: 1 != 0]. Ignoring dependency.
> WARNING: CPU: 4 PID: 0 at arch/x86/kernel/smpboot.c:415 topology_sane.isra.0+0x6b/0x80

To work around this issue, do not use AMD EPYC CPUs for NUMA node configurations.

Bugzilla:2176010

NFS failure during VM migration causes migration failure and source VM coredump

Currently, if the NFS service or server is shut down during virtual machine (VM) migration, the source VM's QEMU is unable to reconnect to the NFS server when it starts running again. As a result, the migration fails and a coredump is initiated on the source VM. Currently, there is no workaround available.

Bugzilla:2058982

PCIe ATS devices do not work on Windows VMs

When you configure a PCIe Address Translation Services (ATS) device in the XML configuration of virtual machine (VM) with a Windows guest operating system, the guest does not enable the ATS device after booting the VM. This is because Windows currently does not support ATS on **virtio** devices.

For more information, see the Red Hat KnowledgeBase .

Bugzilla:2073872

virsh blkiotune --weight command fails to set the correct cgroup I/O controller value

Currently, using the **virsh blkiotune --weight** command to set the VM weight does not work as expected. The command fails to set the correct **io.bfq.weight** value in the cgroup I/O controller interface file. There is no workaround at this time.

Bugzilla:1970830

Starting a VM with an NVIDIA A16 GPU sometimes causes the host GPU to stop working

Currently, if you start a VM that uses an NVIDIA A16 GPU passthrough device, the NVIDIA A16 GPU physical device on the host system in some cases stops working.

To work around the problem, reboot the hypervisor and set the **reset_method** for the GPU device to **bus**:

```
# echo bus > /sys/bus/pci/devices/<DEVICE-PCI-ADDRESS>/reset_method
# cat /sys/bus/pci/devices/<DEVICE-PCI-ADDRESS>/reset_method
bus
```

For details, see the Red Hat Knowledgebase .

Jira:RHEL-7212[1]

Windows VMs might become unresponsive due to storage errors

On virtual machines (VMs) that use Windows guest operating systems, the system in some cases becomes unresponsive when under high I/O load. When this happens, the system logs a **viostor Reset to device, \Device\RaidPort3, was issued** error.

Jira:RHEL-1609[1]

Windows 10 VMs with certain PCI devices might become unresponsive on boot

Currently, a virtual machine (VM) that uses a Windows 10 guest operating system might become unresponsive during boot if a **virtio-win-scsi** PCI device with a local disk back end is attached to the VM. To work around the problem, boot the VM with the **multi_queue** option enabled.

Jira:RHEL-1084[1]

The repair function of virtio-win-guest-tool for the virtio-win drivers does not work

Currently, when using the **Repair** button of **virtio-win-guest-tool** for a **virtio-win** driver, such as the Virtio Balloon Driver, the button has no effect. As a consequence, the driver cannot be reinstalled after being removed on the guest.

Jira:RHEL-1517[1]

Windows 11 VMs with a memory balloon device set might close unexpectedly during reboot

Currently, rebooting virtual machines (VMs) that use a Windows 11 guest operating system and a memory balloon device in some cases fails with a **DRIVER POWER STAT FAILURE** blue-screen error.

Jira:RHEL-935[1]

Resuming a postcopy VM migration fails in some cases

Currently, when performing a postcopy migration of a virtual machine (VM), if a proxy network failure occurs during the RECOVER phase of the migration, the VM becomes unresponsive and the migration cannot be resumed. Instead, the recovery command displays the following error:

> error: Requested operation is not valid: QEMU reports migration is still running

Jira:RHEL-7115

The virtio balloon driver sometimes does not work on Windows 10 VMs

Under certain circumstances, the virtio-balloon driver does not work correctly on virtual machines (VMs) that use a Windows 10 guest operating system. As a consequence, such VMs might not use their assigned memory efficiently.

Jira:RHEL-12118

The virtio file system has suboptimal performance in Windows VMs

Currently, when a virtio file system (virtiofs) is configured on a virtual machine (VM) that uses a Windows guest operating system, the performance of virtiofs in the VM is significantly worse than in VMs that use Linux guests.

Jira:RHEL-1212[1]

Hot-unplugging a storage device on Windows VMs might fail

On virtual machines (VMs) that use a Windows guest operating system, removing a storage device when the VM is running (also known as a device hot-unplug) in some cases fails. As a consequence, the storage device remains attached to the VM and the disk manager service might become unresponsive.

Jira:RHEL-869

Hot plugging CPUs to a Windows VM might cause a system failure

When hot plugging the maximum number of CPUs to a Windows virtual machine (VM) with huge pages enabled, the guest operating system might crash with the following *Stop error*:

> PROCESSOR_START_TIMEOUT

Jira:RHEL-1220

Updating virtio drivers on Windows VMs might fail

When updating the KVM paravirtualized (**virtio**) drivers on a Windows virtual machine (VM), the update might cause the mouse to stop working and the newly installed drivers might not be signed. This problem occurs when updating the **virtio** drivers by installing from the **virtio-win-guest-tools** package, which is a part of the **virtio-win.iso** file.

To work around this problem, update the **virtio** drivers by using Windows Device Manager.

Jira:RHEL-574[1]

TX queue size cannot be changed in VMs that use vhost-kernel

Currently, you cannot set up TX queue size on KVM virtual machines (VMs) that use **vhost-kernel** as a back end for the **virtio** network driver. As a consequence, you can use only the default value of 256 for the TX queue, which might prevent you from optimizing your VM network throughput.

Jira:RHEL-1138[1]

Virtual machines incorrectly report an AMD SRSO vulnerability

RHEL 9.4 virtual machines (VMs) running on a RHEL 9 host with the AMD Zen 3 and 4 CPU architecture incorrectly report a vulnerability to a Speculative Return Stack Overflow (SRSO) attack:

> # lscpu | grep rstack
>
> Vulnerability Spec rstack overflow: Vulnerable: Safe RET, no microcode

The problem is caused by a missing **cpuid** flag and the vulnerability is in fact fully mitigated in VMs under the following conditions:

- You have the updated **linux-firmware** package on the host as described here: cve-2023-20569.

- The host kernel has the mitigation enabled, which is the default behavior. If the mitigation is enabled, **Safe RET** is displayed in the **lscpu** command output on the host.

Jira:RHEL-26152[1]

Virtual machines with a large amount of vCPUs and virtual disks might fail

Currently, assigning a large amount of vCPUs and virtual disks to a RHEL virtual machine (VM) might cause the VM to fail to boot.

To work around this problem, use Small Computer System Interface (SCSI) virtual storage devices instead of block devices if possible. For more details, see: Creating SCSI-based storage pools with vHBA devices by using the CLI

If you need to use virtual block devices, you can also try to reduce the number of interrupt vectors by starting the VM with a **-global virtio-blk-pci.vectors=_<number-of-vectors>_** QEMU option. Try to find a sufficiently low number of interrupt vectors that allows the VM to boot successfully.

Jira:RHEL-32990[1]

Link status shows up on VM, even when status is down of e1000e or igb model interface

Before booting the VM, set the status of Ethernet link **down** for the **e1000** or **igb** model network interface. Despite this, after the VM boots, the network interface keeps the **up** status, because when you set the status of Ethernet link **down** and then stop and re-start the VM, it is automatically set back to **up**. Consequently, the correct state of network interface is not maintained. As a workaround, set the network interface status to **down** inside the VM by using command:

> # ip link set dev eth0 down

Alternatively, you can try to remove and add this network interface again while the VM is running.

Jira:RHEL-21867

Using NBD to migrate a VM storage over a TLS connection does not work correctly

Currently, when migrating a virtual machine (VM) and its storage device by using the Network Block Device (NBD) protocol over a TLS connection, a data race in the TLS handshake might make the migration appear to be successful. However, it causes the QEMU process on the destination VM to become unresponsive to further interactions.

If you can trust your network, you can work around this problem by using plaintext rather than TLS connections for the NBD protocol, which is used during the VM storage migration.

Jira:RHEL-33440

Kdump fails on virtual machines with AMD SEV-SNP

Currently, kdump fails on RHEL 9 virtual machines (VMs) that use the AMD Secure Encrypted Virtualization (SEV) with the Secure Nested Paging (SNP) feature.

Jira:RHEL-10019[1]

11.17. RHEL IN CLOUD ENVIRONMENTS

Cloning or restoring RHEL 9 virtual machines that use LVM on Nutanix AHV causes non-root partitions to disappear

When running a RHEL 9 guest operating system on a virtual machine (VM) hosted on the Nutanix AHV hypervisor, restoring the VM from a snapshot or cloning the VM currently causes non-root partitions in the VM to disappear if the guest is using Logical Volume Management (LVM). As a consequence, the following problems occur:

- After restoring the VM from a snapshot, the VM cannot boot, and instead enters emergency mode.

- A VM created by cloning cannot boot, and instead enters emergency mode.

To work around these problems, do the following in emergency mode of the VM:

1. Remove the LVM system devices file: **rm /etc/lvm/devices/system.devices**

2. Re-create LVM device settings: **vgimportdevices -a**

3. Reboot the VM

This makes it possible for the cloned or restored VM to boot up correctly.

Alternatively, to prevent the issue from occurring, do the following before cloning a VM or creating a VM snapshot:

1. Uncomment the **use_devicesfile = 0** line in the **/etc/lvm/lvm.conf** file

2. Reboot the VM

Bugzilla:2059545[1]

Customizing RHEL 9 guests on ESXi sometimes causes networking problems

Currently, customizing a RHEL 9 guest operating system in the VMware ESXi hypervisor does not work correctly with NetworkManager key files. As a consequence, if the guest is using such a key file, it will have incorrect network settings, such as the IP address or the gateway.

For details and workaround instructions, see the VMware Knowledge Base.

Bugzilla:2037657[1]

RHEL instances on Azure fail to boot if provisioned by cloud-init and configured with an NFSv3 mount entry

Currently, booting a RHEL virtual machine (VM) on the Microsoft Azure cloud platform fails if the VM was provisioned by the **cloud-init** tool and the guest operating system of the VM has an NFSv3 mount entry in the **/etc/fstab** file.

Bugzilla:2081114[1]

Setting static IP in a RHEL virtual machine on a VMware host does not work

Currently, when using RHEL as a guest operating system of a virtual machine (VM) on a VMware host, the DatasourceOVF function does not work correctly. As a consequence, if you use the **cloud-init** utility to set the VM's network to static IP and then reboot the VM, the VM's network will be changed to DHCP.

To work around this issue, see the VMware Knowledge Base.

Jira:RHEL-12122

Large VMs might fail to boot into the debug kernel when the kmemleak option is enabled

When attempting to boot a RHEL 9 virtual machine (VM) into the debug kernel, the booting might fail with the following error if the machine kernel is using the **kmemleak=on** argument.

> Cannot open access to console, the root account is locked.
> See sulogin(8) man page for more details.
>
> Press Enter to continue.

This problem affects mainly large VMs because they spend more time in the boot sequence.

To work around the problem, edit the **/etc/fstab** file on the machine and add extra timeout options to the **/boot** and **/boot/efi** mount points. For example:

> UUID=e43ead51-b364-419e-92fc-b1f363f19e49 /boot xfs defaults,x-systemd.device-timeout=600,x-systemd.mount-timeout=600 0 0
>
> UUID=7B77-95E7 /boot/efi vfat defaults,uid=0,gid=0,umask=077,shortname=winnt,x-systemd.device-timeout=600,x-systemd.mount-timeout=600 0 2

Jira:RHELDOCS-16979[1]

11.18. SUPPORTABILITY

Timeout when running sos report on IBM Power Systems, Little Endian

When running the **sos report** command on IBM Power Systems, Little Endian with hundreds or

thousands of CPUs, the processor plugin reaches its default timeout of 300 seconds when collecting huge content of the **/sys/devices/system/cpu** directory. As a workaround, increase the plugin's timeout accordingly:

- For one-time setting, run:

```
# sos report -k processor.timeout=1800
```

- For a permanent change, edit the **[plugin_options]** section of the **/etc/sos/sos.conf** file:

```
[plugin_options]
# Specify any plugin options and their values here. These options take the form
# plugin_name.option_name = value
#rpm.rpmva = off
processor.timeout = 1800
```

The example value is set to 1800. The particular timeout value highly depends on a specific system. To set the plugin's timeout appropriately, you can first estimate the time needed to collect the one plugin with no timeout by running the following command:

```
# time sos report -o processor -k processor.timeout=0 --batch --build
```

Bugzilla:1869561[1]

11.19. CONTAINERS

Running systemd within an older container image does not work

Running systemd within an older container image, for example, **centos:7**, does not work:

```
$ podman run --rm -ti centos:7 /usr/lib/systemd/systemd
 Storing signatures
 Failed to mount cgroup at /sys/fs/cgroup/systemd: Operation not permitted
 [!!!!!!] Failed to mount API filesystems, freezing.
```

To work around this problem, use the following commands:

```
# mkdir /sys/fs/cgroup/systemd
# mount none -t cgroup -o none,name=systemd /sys/fs/cgroup/systemd
# podman run --runtime /usr/bin/crun --annotation=run.oci.systemd.force_cgroup_v1=/sys/fs/cgroup -
-rm -ti centos:7 /usr/lib/systemd/systemd
```

Jira:RHELPLAN-96940[1]

Root filesystem are not expanded by default

When you use a base container image, that does not include **cloud-init** to create an AMI or QCOW2 container image by using **bootc-image-builder**, the root filesystem size is not expanded dynamically on boot to the full size of the provisioned virtual disk.

To workaround this issue, apply one of the following available options:

- Include **cloud-init** in the image.

- Include custom logic in the container image to expand the root filesystem, for example:

```
/usr/bin/growpart /dev/vda 4
unshare -m bin/sh -c 'mount -o remount,rw /sysroot && xfs_growfs /sysroot'
```

- Include a custom logic to use the additional space for secondary filesystems, for example, **/var/lib/containers**.

NOTE

By default, the physical root storage is mounted at the **/sysroot** partition.

Jira:RHEL-33208

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-15907, Jira:RHEL-5142, Jira:RHEL-5133, Jira:RHEL-5130, Jira:RHEL-16830, Jira:RHEL-17175, Jira:RHEL-5111
NetworkManager	Jira:RHEL-1441, Jira:RHEL-1471, Jira:RHEL-16470, Jira:RHEL-1469, Jira:RHEL-24337, Jira:RHEL-5852, Bugzilla:1894877, Jira:RHEL-17619
Release Notes	Jira:RHELDOCS-17841, Jira:RHELDOCS-16861, Jira:RHELDOCS-16760, Jira:RHELDOCS-17520, Jira:RHELDOCS-17803, Jira:RHELDOCS-16756, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102, Jira:RHELDOCS-17166, Jira:RHELDOCS-17309, Jira:RHELDOCS-17545, Jira:RHELDOCS-17518, Jira:RHELDOCS-17989, Jira:RHELDOCS-17702, Jira:RHELDOCS-17917, Jira:RHELDOCS-16979
anaconda	Jira:RHEL-11384, Jira:RHEL-13150, Jira:RHEL-4766, Jira:RHEL-5638, Jira:RHEL-10216, Jira:RHEL-2250, Jira:RHEL-17205, Bugzilla:2127473, Bugzilla:2050140, Bugzilla:1877697, Jira:RHEL-4707, Jira:RHEL-4711, Bugzilla:1997832, Jira:RHEL-4741, Bugzilla:2115783, Jira:RHEL-4762, Jira:RHEL-4737, Jira:RHEL-9633
ansible-collection-microsoft-sql	Jira:RHEL-16342, Jira:RHEL-19092, Jira:RHEL-19091, Jira:RHEL-3540
ansible-freeipa	Jira:RHEL-4962, Jira:RHEL-16934, Jira:RHEL-16939, Jira:RHEL-19134, Jira:RHEL-19130
audit	Jira:RHEL-14896
bacula	Jira:RHEL-6856
bcc	Jira:RHEL-16325
bind	Bugzilla:1984982
boom-boot	Jira:RHEL-16813
bootc-image-builder-container	Jira:RHEL-34054
certmonger	Jira:RHEL-22302
chrony	Jira:RHEL-6522

Component	Tickets
clang	Jira:RHEL-9346
cloud-init	Jira:RHEL-7278, Jira:RHEL-7311, Jira:RHEL-12122
cmake	Jira:RHEL-7393
cockpit-appstream	Bugzilla:2030836
cockpit-machines	Jira:RHEL-17434, Bugzilla:2173584, Jira:RHEL-31993
crash	Jira:RHEL-9009
createrepo_c	Bugzilla:2056318
crypto-policies	Jira:RHEL-15925, Jira:RHEL-2735
cyrus-sasl	Bugzilla:1995600
device-mapper-multipath	Jira:RHEL-6678, Jira:RHEL-1729, Jira:RHEL-4998, Jira:RHEL-986, Jira:RHEL-1830, Jira:RHEL-17234, Bugzilla:2033080, Bugzilla:2011699, Bugzilla:1926147
distribution	Jira:RHEL-17089, Jira:RHEL-6973, Jira:RHEL-18157, Jira:RHEL-22385
dnf	Bugzilla:2073510
dnf-plugins-core	Jira:RHEL-4600
edk2	Bugzilla:1935497
elfutils	Jira:RHEL-12489
fapolicyd	Bugzilla:2054740, Jira:RHEL-24345, Jira:RHEL-520
firewalld	Jira:RHEL-427, Jira:RHEL-14485, Jira:RHEL-17708
gcc	Jira:RHEL-17638
gcc-toolset-13-binutils	Jira:RHEL-23798
gcc-toolset-13-gcc	Jira:RHEL-16998
gimp	Bugzilla:2047161
git	Jira:RHEL-17100

Component	Tickets
git-lfs	Jira:RHEL-17101
glibc	Jira:RHEL-14383, Jira:RHEL-17157, Jira:RHEL-2491, Jira:RHEL-19862, Jira:RHEL-16643, Jira:RHEL-12362, Jira:RHEL-3397, Jira:RHEL-2123
gnupg2	Bugzilla:2070722
gnutls	Jira:RHEL-14891, Bugzilla:2108532
golang	Jira:RHEL-11871, Bugzilla:2111072, Bugzilla:2092016
grafana	Jira:RHEL-7505
grub2	Jira:RHEL-10288
gtk3	Jira:RHEL-11924
httpd	Jira:RHEL-6600
ipa	Jira:RHEL-11652, Jira:RHEL-23377, Bugzilla:1513934, Jira:RHEL-9984, Jira:RHEL-22313, Jira:RHEL-12143, Bugzilla:2084180, Bugzilla:2094673, Bugzilla:2057471, Jira:RHEL-12154, Jira:RHEL-4955
iptables	Jira:RHEL-14147
jmc-core	Bugzilla:1980981
kdump-anaconda-addon	Jira:RHEL-11196
kernel	Jira:RHEL-11597, Bugzilla:2041883, Bugzilla:1613522, Bugzilla:1995338, Bugzilla:1570255, Bugzilla:2177256, Bugzilla:2178699, Bugzilla:2023416, Bugzilla:2021672, Bugzilla:2027304, Bugzilla:1660337, Bugzilla:1955275, Bugzilla:2142102, Bugzilla:2040643, Bugzilla:2186375, Bugzilla:2183538, Bugzilla:2206599, Bugzilla:2167783, Bugzilla:2000616, Bugzilla:2013650, Bugzilla:2132480, Bugzilla:2059545, Bugzilla:2005173, Bugzilla:2128610, Bugzilla:2129288, Bugzilla:2013884, Bugzilla:2149989
kernel / BPF	Jira:RHEL-10691
kernel / Crypto	Jira:RHEL-20145
kernel / DMA Engine	Jira:RHEL-10097
kernel / Debugging-Tracing / rtla	Jira:RHEL-10079

Component	Tickets
kernel / Kernel-Core	Jira:RHEL-25967
kernel / Networking / IPSec	Jira:RHEL-1015
kernel / Networking / NIC Drivers	Jira:RHEL-9308, Jira:RHEL-24618, Jira:RHEL-9897
kernel / Networking / Netfilter	Jira:RHEL-16630
kernel / Networking / Protocol / tcp	Jira:RHEL-21223, Jira:RHEL-5736
kernel / Platform Enablement / NVMe	Jira:RHEL-21545, Jira:RHEL-14751, Jira:RHEL-8171, Jira:RHEL-8164
kernel / Platform Enablement / ppc64	Jira:RHEL-15404
kernel / Security / TPM	Jira:RHEL-18985
kernel / Storage / Device Mapper / Crypt	Jira:RHEL-23572
kernel / Storage / Multiple Devices (MD)	Jira:RHEL-30730
kernel / Storage / Storage Drivers	Jira:RHEL-8466, Jira:RHEL-8104, Jira:RHEL-25730
kernel / Virtualization	Jira:RHEL-1138
kernel / Virtualization / KVM	Jira:RHEL-2815, Jira:RHEL-13007, Jira:RHEL-7212, Jira:RHEL-26152, Jira:RHEL-10019
kernel-rt	Bugzilla:2181571
kernel-rt / Other	Jira:RHEL-9318
kexec-tools	Bugzilla:2113873, Bugzilla:2064708
keylime	Jira:RHEL-11867, Jira:RHEL-1518
kmod	Bugzilla:2103605
kmod-kvdo	Jira:RHEL-8354

Component	Tickets
krb5	Jira:RHEL-4902, Bugzilla:2060798, Jira:RHEL-4875, Jira:RHEL-4889, Bugzilla:2016312, Jira:RHEL-4888
libabigail	Jira:RHEL-16629
libdnf	Jira:RHEL-11238
libkcapi	Jira:RHEL-15298, Jira:RHEL-5367
libotr	Bugzilla:2086562
librepo	Jira:RHEL-11240
libreswan	Jira:RHEL-12278
librhsm	Jira:RHEL-14224
libsepol	Jira:RHEL-16233
libvirt	Bugzilla:2143158, Bugzilla:2078693
libvirt / General	Jira:RHEL-7568, Jira:RHEL-7043
libvirt / Storage	Jira:RHEL-7528, Jira:RHEL-7416
libxcrypt	Bugzilla:2034569
libzip	Jira:RHEL-17567
linuxptp	Jira:RHEL-2026
llvm-toolset	Jira:RHEL-9283
lvm2	Jira:RHEL-8357, Bugzilla:2038183
make	Jira:RHEL-22829
mariadb	Jira:RHEL-3638
maven	Jira:RHEL-13046
mysql	Bugzilla:1991500
nettle	Jira:RHEL-14890

Component	Tickets
nfs-utils	Bugzilla:2081114
nftables	Jira:RHEL-5980, Jira:RHEL-14191
nginx	Jira:RHEL-14713
nmstate	Jira:RHEL-1434, Jira:RHEL-1605, Jira:RHEL-1438, Jira:RHEL-19142, Jira:RHEL-1420, Jira:RHEL-1425
nvme-cli	Jira:RHEL-1492
nvme-stas	Bugzilla:1893841
open-vm-tools	Bugzilla:2037657
opencryptoki	Jira:RHEL-11412
opensc	Jira:RHEL-4079
openscap	Bugzilla:2161499
openslp	Jira:RHEL-6995
openssh	Jira:RHEL-5222, Jira:RHEL-2469, Bugzilla:2056884
openssl	Jira:RHEL-23474, Jira:RHEL-17193, Bugzilla:2168665, Bugzilla:1975836, Bugzilla:1681178, Bugzilla:1685470
osbuild	Jira:RHEL-4655
osbuild-composer	Bugzilla:2173928, Jira:RHEL-7999, Jira:RHEL-4649
oscap-anaconda-addon	Jira:RHEL-1824, Jira:RHELPLAN-44202
p11-kit	Jira:RHEL-14834
papi	Jira:RHEL-9333
pause-container	Bugzilla:2106816
pcp	Jira:RHEL-2317
pcs	Jira:RHEL-7672, Jira:RHEL-7582, Jira:RHEL-7724, Jira:RHEL-7746, Jira:RHEL-7744, Jira:RHEL-7669, Jira:RHEL-7730

Component	Tickets
php	Jira:RHEL-14699
pki-core	Bugzilla:2084181
podman	Jira:RHELPLAN-167829, Jira:RHELPLAN-167796, Jira:RHELPLAN-167823, Jira:RHELPLAN-168180, Jira:RHELPLAN-168185, Jira:RHELPLAN-154436, Bugzilla:2069279
policycoreutils	Jira:RHEL-24462, Jira:RHEL-25263
postgresql	Jira:RHEL-3635
procps-ng	Jira:RHEL-16278
python3.11-lxml	Bugzilla:2157708
qemu-kvm	Jira:RHEL-11597, Jira:RHEL-16695, Bugzilla:1965079, Bugzilla:1951814, Bugzilla:2060839, Bugzilla:2014229, Jira:RHEL-7335, Jira:RHEL-7336, Bugzilla:1915715, Jira:RHEL-13335, Jira:RHEL-333, Bugzilla:2176010, Bugzilla:2058982, Bugzilla:2073872, Jira:RHEL-7478, Jira:RHEL-32990
qemu-kvm / Devices	Jira:RHEL-1220
qemu-kvm / Graphics	Jira:RHEL-7135
qemu-kvm / Live Migration	Jira:RHEL-13004, Jira:RHEL-7096, Jira:RHEL-7115
qemu-kvm / Networking	Jira:RHEL-7337, Jira:RHEL-21867
qemu-kvm / Storage / NBD	Jira:RHEL-33440
realtime-tests	Jira:RHEL-9910
rear	Jira:RHEL-16864, Jira:RHEL-10478, Jira:RHEL-6984, Jira:RHEL-17393, Jira:RHEL-24847
restore	Bugzilla:1997366
rhel-bootc-container	Jira:RHEL-33208

Component	Tickets
rhel-system-roles	Jira:RHEL-1535, Jira:RHEL-16976, Jira:RHEL-16542, Jira:RHEL-16552, Jira:RHEL-19579, Jira:RHEL-17668, Jira:RHEL-21133, Jira:RHEL-16964, Jira:RHEL-16541, Jira:RHEL-15932, Jira:RHEL-15439, Jira:RHEL-18962, Jira:RHEL-16212, Jira:RHEL-15876, Jira:RHEL-21117, Jira:RHEL-16974, Jira:RHEL-5972, Jira:RHEL-15037, Jira:RHEL-19046, Jira:RHEL-18026, Jira:RHEL-1683, Jira:RHEL-3353, Jira:RHEL-19043, Jira:RHEL-19040, Jira:RHEL-17875, Jira:RHEL-5274, Jira:RHEL-15870, Jira:RHEL-15909, Jira:RHEL-22228, Jira:RHEL-25508, Jira:RHEL-21401, Jira:RHEL-22309, Bugzilla:1999770, Bugzilla:2123859, Jira:RHEL-1172, Bugzilla:2186218
rsyslog	Jira:RHEL-943, Jira:RHEL-937, Jira:RHEL-5196
rteval	Jira:RHEL-9912
rust	Jira:RHEL-12963
s390utils	Bugzilla:1932480
samba	Jira:RHEL-16476
scap-security-guide	Jira:RHEL-21425, Jira:RHEL-1800, Bugzilla:2038978
selinux-policy	Jira:RHEL-12591, Jira:RHEL-21452, Jira:RHEL-1548, Jira:RHEL-18219, Jira:RHEL-14246, Jira:RHEL-1551, Jira:RHEL-11174, Jira:RHEL-10087, Jira:RHEL-1553, Jira:RHEL-14289, Jira:RHEL-5032, Jira:RHEL-15432, Jira:RHEL-19051, Jira:RHEL-11792, Bugzilla:2064274, Jira:RHEL-28814
sos	Bugzilla:1869561
sssd	Jira:SSSD-7015, Bugzilla:1608496
sssd_kcm	Jira:SSSD-7015
stratis-cli	Jira:RHEL-2265
stratisd	Jira:RHEL-12898, Jira:RHEL-16736
stunnel	Jira:RHEL-2468
subscription-manager	Bugzilla:2163716, Bugzilla:2136694
synce4l	Jira:RHEL-10089
sysstat	Jira:RHEL-12009, Jira:RHEL-26275

Component	Tickets
systemd	Bugzilla:2018112, Jira:RHEL-6105
systemtap	Jira:RHEL-12488
tigervnc	Bugzilla:2060308
tuna	Jira:RHEL-8859
tuned	Bugzilla:2113900
udisks2	Bugzilla:2213769
unbound	Bugzilla:2070495
vdo	Jira:RHEL-30525
virt-v2v	Bugzilla:2168082
virtio-win	Jira:RHEL-11810, Jira:RHEL-11366, Jira:RHEL-1609, Jira:RHEL-869
virtio-win / distribution	Jira:RHEL-1517, Jira:RHEL-1860, Jira:RHEL-574
virtio-win / virtio-win-prewhql	Jira:RHEL-1084, Jira:RHEL-935, Jira:RHEL-12118, Jira:RHEL-1212
webkit2gtk3	Jira:RHEL-4157
xdp-tools	Jira:RHEL-3382

Component	Tickets
other	Jira:RHELDOCS-17369, Jira:RHELDOCS-17263, Jira:RHELDOCS-17060, Jira:RHELDOCS-17056, Jira:RHELDOCS-16721, Jira:RHELDOCS-17372, Jira:RHELDOCS-16970, Jira:RHELPLAN-169666, Jira:RHELDOCS-17000, Jira:RHELDOCS-16241, Jira:RHELDOCS-16955, Jira:RHELDOCS-17053, Jira:RHELDOCS-17792, Jira:SSSD-6184, Jira:RHELDOCS-17040, Bugzilla:2020529, Jira:RHELPLAN-27394, Jira:RHELPLAN-27737, Jira:RHELDOCS-16861, Jira:RHELDOCS-17050, Jira:RHELDOCS-17520, Jira:RHELDOCS-17752, Jira:RHELDOCS-17803, Jira:RHELDOCS-17468, Jira:RHELDOCS-17733, Bugzilla:1927780, Jira:RHELPLAN-110763, Bugzilla:1935544, Bugzilla:2089200, Jira:RHELPLAN-99136, Jira:RHELPLAN-103232, Bugzilla:1899167, Bugzilla:1979521, Jira:RHELPLAN-100087, Jira:RHELPLAN-100639, Bugzilla:2058153, Jira:RHELPLAN-113995, Jira:RHELPLAN-98983, Jira:RHELPLAN-131882, Jira:RHELPLAN-139805, Jira:RHELDOCS-16756, Jira:RHELPLAN-153267, Jira:RHELDOCS-16300, Jira:RHELDOCS-16432, Jira:RHELDOCS-16393, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102, Jira:RHELDOCS-17015, Jira:RHELDOCS-18049, Jira:RHELDOCS-17135, Jira:RHELDOCS-17545, Jira:RHELDOCS-17038, Jira:RHELDOCS-17495, Jira:RHELDOCS-17518, Jira:RHELDOCS-17462, Jira:RHELDOCS-18106, Jira:RHELPLAN-157225, Bugzilla:1640697, Bugzilla:1697896, Bugzilla:2047713, Jira:RHELPLAN-96940, Jira:RHELPLAN-117234, Jira:RHELPLAN-119001, Jira:RHELPLAN-119852, Bugzilla:2077767, Bugzilla:2053598, Bugzilla:2082303, Jira:RHELPLAN-121049, Jira:RHELPLAN-157939, Jira:RHELPLAN-109613, Bugzilla:2160619, Jira:RHELDOCS-18064, Jira:RHELDOCS-16427, Bugzilla:2173992, Bugzilla:2185048, Bugzilla:1970830, Jira:RHELDOCS-16574, Jira:RHELDOCS-17719, Jira:RHELDOCS-17720

APPENDIX B. REVISION HISTORY

0.0-1

Wed May 01 2024, Gabriela Fialová (gfialova@redhat.com)

- Release of the Red Hat Enterprise Linux 9.4 Release Notes.

0.0-0

Wed March 27 2024, Gabriela Fialová (gfialova@redhat.com)

- Release of the Red Hat Enterprise Linux 9.4 Beta Release Notes.