Red Hat Enterprise Linux 8.9

8.9 Release Notes

Release Notes for Red Hat Enterprise Linux 8.9

Red Hat Enterprise Linux 8.9 8.9 Release Notes

Release Notes for Red Hat Enterprise Linux 8.9

Legal Notice

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 8.9 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 8.9

Installer and image creation
Key highlights for image builder:

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

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

Security
Key security-related highlights:

- **OpenSCAP** was rebased to version 1.3.8.

- **ANSSI-BP-028** SCAP security profiles were updated to version 2.0.

- **SCAP Security Guide** now contains improved rules to provide more consistent interactive user configuration and the DISA STIG profile supports **audit_rules_login_events_faillock**.

See New features – Security for more information.

Dynamic programming languages, web and database servers
Node.js 20 is now available as a new module stream.

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 8.9:

- **Valgrind 3.21**

- **SystemTap 4.9**

- **elfutils 0.189**

Updated performance monitoring tools
The following performance monitoring tools have been updated in RHEL 8.9:

- **Grafana 9.2.10**

- **grafana-pcp 5.1.1**

Updated compiler toolsets
The following compiler toolsets have been updated in RHEL 8.9:

- **GCC Toolset 13** (new)

- **LLVM Toolset 16.0.6**

- **Rust Toolset 1.71.1**

- **Go Toolset 1.20.10**

See New features – Compilers and development tools for more information.

Java implementations in RHEL 8
The RHEL 8 AppStream repository includes:

- The **java-21-openjdk** packages, which provide the OpenJDK 21 Java Runtime Environment and the OpenJDK 21 Java Software Development Kit.

- The **java-17-openjdk** packages, which provide the OpenJDK 17 Java Runtime Environment and the OpenJDK 17 Java Software Development Kit.

- The **java-11-openjdk** packages, which provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Software Development Kit.

- The **java-1.8.0-openjdk** packages, which provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Software Development Kit.

The Red Hat build of OpenJDK packages share a single set of binaries between its portable Linux releases, RHEL 8.9 and later releases. Because of this update, there is a change in the process of rebuilding the OpenJDK packages on RHEL from the source RPM. For more information about the new rebuilding process, see the **README.md** file which is available in the SRPM package of the Red Hat build of OpenJDK and is also installed by the **java-*-openjdk-headless** packages under the **/usr/share/doc** tree.

For more information, see OpenJDK documentation.

1.2. IN-PLACE UPGRADE AND OS CONVERSION

In-place upgrade from RHEL 7 to RHEL 8
The possible in-place upgrade paths currently are:

- From RHEL 7.9 to RHEL 8.6 RHEL 8.8 and RHEL 8.9 on the 64-bit Intel, IBM POWER 8 (little endian), and IBM Z architectures

- From RHEL 7.9 to RHEL 8.6 and RHEL 8.8 on systems with SAP HANA on the 64-bit Intel architecture.

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 7 to RHEL 8 .

For instructions on performing an in-place upgrade on systems with SAP environments, see How to in-place upgrade SAP environments from RHEL 7 to RHEL 8.

Notable enhancements include:

- Requirements on disk space have been significantly reduced on systems with XFS filesystems formatted with **ftype=0**.

- Disk images created during the upgrade process for upgrade purposes now have dynamic sizes. The **LEAPP_OVL_SIZE** environment variable is not needed anymore.

- Issues with the calculation of the required free space on existing disk partitions have been fixed. The missing free disk space is now correctly detected before the required reboot of the system, and the report correctly displays file systems that do not have enough free space to proceed the upgrade RPM transaction.

- Third-party drivers can now be managed during the in-place upgrade process using custom leapp actors.

- An overview of the pre-upgrade and upgrade reports is now printed in the terminal.

- Upgrades of RHEL Real Time and RHEL Real Time for Network Functions Virtualization (NFV) in Red Hat OpenStack Platform are now supported.

In-place upgrade from RHEL 6 to RHEL 8
It is not possible to perform an in-place upgrade directly from RHEL 6 to RHEL 8. However, you can perform an in-place upgrade from RHEL 6 to RHEL 7 and then perform a second in-place upgrade to RHEL 8. For more information, see Upgrading from RHEL 6 to RHEL 7 .

In-place upgrade from RHEL 8 to RHEL 9
Instructions on how to perform an in-place upgrade from RHEL 8 to RHEL 9 using the Leapp utility are provided by the document Upgrading from RHEL 8 to RHEL 9 . Major differences between RHEL 8 and RHEL 9 are documented in Considerations in adopting RHEL 9 .

Conversion from a different Linux distribution to RHEL
If you are using Alma Linux 8, CentOS Linux 8, Oracle Linux 8, or Rocky Linux 8, you can convert your operating system to RHEL 8 using the Red Hat-supported **Convert2RHEL** utility. For more information, see Converting from an RPM-based Linux distribution to RHEL .

If you are using an earlier version of CentOS Linux or Oracle Linux, namely versions 6 or 7, you can convert your operating system to RHEL and then perform an in-place upgrade to RHEL 8. Note that CentOS Linux 6 and Oracle Linux 6 conversions use the unsupported **Convert2RHEL** utility. For more information on unsupported conversions, see How to perform an unsupported conversion from a RHEL-derived Linux distribution to RHEL.

For information regarding how Red Hat supports conversions from other Linux distributions to RHEL, see the Convert2RHEL Support Policy document .

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

- Product Life Cycle Checker

- Kickstart Generator

- Kickstart Converter

- Red Hat Enterprise Linux Upgrade Helper

- Red Hat Satellite Upgrade Helper

- Red Hat Code Browser

- JVM Options Configuration Tool

- Red Hat CVE Checker

- Red Hat Product Certificates

- Load Balancer Configuration Tool

- Yum Repository Configuration Helper

- Red Hat Memory Analyzer

- Kernel Oops Analyzer

- Red Hat Product Errata Advisory Checker

1.4. ADDITIONAL RESOURCES

- **Capabilities and limits** of Red Hat Enterprise Linux 8 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 8.

- Major **differences between RHEL 7 and RHEL 8**, including removed functionality, are documented in Considerations in adopting RHEL 8 .

- Instructions on how to perform an **in-place upgrade from RHEL 7 to RHEL 8** are provided by the document Upgrading from RHEL 7 to RHEL 8 .

- The **Red Hat Insights** service, which enables you to proactively identify, examine, and resolve known technical issues, is now 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.
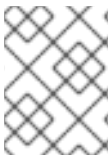
> **NOTE**
>
> Release notes include links to access the original tracking tickets. Private tickets have no links and instead feature this footnote.[1]

---

[1] Release notes include links to access the original tracking tickets. Private tickets have no links and instead feature this footnote.

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 8.9 is distributed with the kernel version 4.18.0-513.5.1, which provides support for the following architectures:

- AMD and Intel 64-bit architectures

- The 64-bit ARM architecture

- IBM Power Systems, Little Endian

- 64-bit IBM Z

Make sure you purchase the appropriate subscription for each architecture. For more information, see Get Started with Red Hat Enterprise Linux - additional architectures . For a list of available subscriptions, see Subscription Utilization on the Customer Portal.

CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 8

3.1. INSTALLATION

Red Hat Enterprise Linux 8 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:

- Binary DVD ISO: A full installation image that contains the BaseOS and AppStream repositories and allows you to complete the installation without additional repositories.

    > **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 Binary DVD ISO image.

See the Performing a standard RHEL 8 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 8 installation document.

3.2. REPOSITORIES

Red Hat Enterprise Linux 8 is distributed through two main repositories:

- BaseOS

- AppStream

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

Content in the BaseOS repository is intended to provide the core set of the underlying OS functionality that provides the foundation for all installations. This content is available in the RPM format and is subject to support terms similar to those in previous releases of RHEL. For a list of packages distributed through BaseOS, see the Package manifest.

Content in the Application Stream repository includes additional user space applications, runtime languages, and databases in support of the varied workloads and use cases. Application Streams are available in the familiar RPM format, as an extension to the RPM format called *modules*, or as Software Collections. For a list of packages available in AppStream, see the Package manifest.

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 8 repositories, see the Package manifest.

3.3. APPLICATION STREAMS

Red Hat Enterprise Linux 8 introduces the concept of Application Streams. Multiple versions of user-space components are now delivered and updated more frequently than the core operating system packages. This provides greater flexibility to customize Red Hat Enterprise Linux without impacting the underlying stability of the platform or specific deployments.

Components made available as Application Streams can be packaged as modules or RPM packages and are delivered through the AppStream repository in RHEL 8. Each Application Stream component has a given life cycle, either the same as RHEL 8 or shorter. For details, see Red Hat Enterprise Linux Life Cycle.

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.

Module streams represent versions of the Application Stream components. For example, several streams (versions) of the PostgreSQL database server are available in the **postgresql** module with the default **postgresql:10** stream. Only one module stream can be installed on the system. Different versions can be used in separate containers.

Detailed module commands are described in the Installing, managing, and removing user-space components document. For a list of modules available in AppStream, see the Package manifest.

3.4. PACKAGE MANAGEMENT WITH YUM/DNF

On Red Hat Enterprise Linux 8, installing software is ensured by the **YUM** tool, which is based on the **DNF** technology. We deliberately adhere to usage of the **yum** term for consistency with previous major versions of RHEL. However, if you type **dnf** instead of **yum**, the command works as expected because **yum** is an alias to **dnf** for compatibility.

For more details, see the following documentation:

- Installing, managing, and removing user-space components

- Considerations in adopting RHEL 8

CHAPTER 4. NEW FEATURES

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 8.9.

4.1. INSTALLER AND IMAGE CREATION

Support to both legacy and UEFI boot for AWS EC2 images

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

Jira:RHELDOCS-16339[1]

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

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

Bugzilla:1770969

New **network** kickstart options to control DNS handling

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

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

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

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

Bugzilla:1656662[1]

4.2. SECURITY

opencryptoki rebased to 3.21.0

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

- Concurrent hardware security module (HSM) master key changes

- The **protected-key** option to transform a chosen key into a protected key

- Additional key types, such as DH, DSA, and generic secret key types

- EP11 host library version 4

- AES-XTS key type

- IBM-specific Kyber key type and mechanism

- Additional IBM-specific Dilithium key round 2 and 3 variants

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

p11sak set-key-attr

To modify keys

p11sak copy-key

To copy keys

p11sak import-key

To import keys

p11sak export-key

To export keys

Bugzilla:2159697[1]

fapolicyd now provides rule numbers for troubleshooting

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

Jira:RHEL-628

ANSSI-BP-028 security profiles updated to version 2.0

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

- ANSSI-BP-028 Minimal Level

- ANSSI-BP-028 Intermediary Level

- ANSSI-BP-028 Enhanced Level

- ANSSI-BP-028 High Level

Bugzilla:2155789

Better definition of interactive users

The rules in the **scap-security-guide** package were improved to provide more consistent interactive user configuration. Previously, some rules used different approaches for identifying interactive and non-interactive users. With this update, we have unified the definitions of interactive users. User accounts with UID greater than or equal to 1000 are now considered interactive, with the exception of the **nobody** and **nfsnobody** accounts and with the exception of accounts that use **/sbin/nologin** as the login shell.

This change affects the following rules:

- **accounts_umask_interactive_users**

- **accounts_user_dot_user_ownership**

- **accounts_user_dot_group_ownership**

- **accounts_user_dot_no_world_writable_programs**

- **accounts_user_interactive_home_directory_defined**

- **accounts_user_interactive_home_directory_exists**

- **accounts_users_home_files_groupownership**

- **accounts_users_home_files_ownership**

- **accounts_users_home_files_permissions**

- **file_groupownership_home_directories**

- **file_ownership_home_directories**

- **file_permissions_home_directories**

- **file_permissions_home_dirs**

- **no_forward_files**

Bugzilla:2157877, Bugzilla:2178740

The DISA STIG profile now supports audit_rules_login_events_faillock

With this enhancement, the SCAP Security Guide **audit_rules_login_events_faillock** rule, which references STIG ID RHEL-08-030590, has been added to the DISA STIG profile for RHEL 8. This rule checks if the Audit daemon is configured to record any attempts to modify login event logs stored in the **/var/log/faillock** directory.

Bugzilla:2167999

OpenSCAP rebased to 1.3.8

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

- Fixed **systemd** probes to not ignore some **systemd** units

- Added offline capabilities to the **shadow** OVAL probe

- Added offline capabilities to the **sysctl** OVAL probe

- Added **auristorfs** to the list of network filesystems

- Created a workaround for issues with tailoring files produced by the **autotailor** utility

Bugzilla:2217441

SCAP Security Guide rebased to version 0.1.69

The SCAP Security Guide (SSG) packages have been rebased to upstream version 0.1.69. This version

provides various enhancements and bug fixes, most notably three new SCAP profiles for RHEL 9 which are aligned with three levels of the CCN-STIC-610A22 Guide issued by the National Cryptologic Center of Spain in 2022-10:

- CCN Red Hat Enterprise Linux 9 - Basic

- CCN Red Hat Enterprise Linux 9 - Intermediate

- CCN Red Hat Enterprise Linux 9 - Advanced

Bugzilla:2221695

FIPS-enabled in-place upgrades from RHEL 8.8 and later to RHEL 9.2 and later are supported

With the release of the RHBA-2023:3824 advisory, you can perform an in-place upgrade of a RHEL 8.8 and later system to a RHEL 9.2 and later system with FIPS mode enabled.

Bugzilla:2097003

crypto-policies permitted_enctypes no longer break replications in FIPS mode

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

With this update, the values of the **permitted_enctypes krb5** configuration option depend on the **mac** and **cipher crypto-policy** values. That allows the prioritization of the interoperable encryption types by default.

As additional results of this update, the **arcfour-hmac-md5** option is available only in the **LEGACY:AD-SUPPORT** subpolicy and the **aes256-cts-hmac-sha1-96** is no longer available in the **FUTURE** policy.

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

Bugzilla:2219912

Audit now supports FANOTIFY record fields

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

- **fan_type** to specify the type of a **FANOTIFY** event

- **fan_info** to specify additional context information

- **sub_trust** and **obj_trust** to indicate trust levels for a subject and an object involved in an event

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

Bugzilla:2216666

New SELinux boolean to allow QEMU Guest Agent executing confined commands

Previously, commands that were supposed to execute in a confined context through the QEMU Guest Agent daemon program, such as **mount**, failed with an Access Vector Cache (AVC) denial. To be able to execute these commands, the **guest-agent** must run in the **virt_qemu_ga_unconfined_t** domain.

Therefore, this update adds the SELinux policy boolean **virt_qemu_ga_run_unconfined** that allows **guest-agent** to make the transition to **virt_qemu_ga_unconfined_t** for executables located in any of the following directories:

- **/etc/qemu-ga/fsfreeze-hook.d/**

- **/usr/libexec/qemu-ga/fsfreeze-hook.d/**

- **/var/run/qemu-ga/fsfreeze-hook.d/**

In addition, the necessary rules for transitions for the **qemu-ga** daemon have been added to the SELinux policy boolean.

As a result, you can now execute confined commands through the QEMU Guest Agent without AVC denials by enabling the **virt_qemu_ga_run_unconfined** boolean.

Bugzilla:2093355

4.3. INFRASTRUCTURE SERVICES

Postfix now supports SRV lookups

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

use_srv_lookup

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

allow_srv_lookup_fallback

You can use a cascading approach to locating a service.

ignore_srv_lookup_error

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

Bugzilla:1787010

You can now specify TLS 1.3 cipher suites in vsftpd

With this enhancement, you can use the new **ssl_ciphersuites** option to configure which cipher suites **vsftpd** uses. As a result, you can specify TLS 1.3 cipher suites that differ from the previous TLS versions. To specify multiple cipher suites, separate entries with colons (:).

Bugzilla:2069733

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

With this enhancement, you can now use the Generic LF-to-CRLF driver, which converts LF characters to CR+LF characters for printers accepting files with CR+LF characters. The carriage return (CR) and line feed (LF) are control characters that mark the end of lines. As a result, by using this driver, you can

send an LF character terminated file from your application to a printer accepting only CR+LF characters. The Generic LF-to-CRLF driver is a renamed version of the **text-only** driver from RHEL 7. The new name reflects its actual functionality.

Bugzilla:2118406[1]

4.4. NETWORKING

iproute rebased to version 6.2.0

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

- The new **ip stats** command manages and shows interface statistics. By default, the **ip stats show** command displays statistics for all network devices, including bridges and bonds. You can filter the output by using the **dev** and **group** options. For further details, see the **ip-stats(8)** man page.

- The **ss** utility now provides the **-T** (**--threads**) option to display thread information, which extends the **-p** (**--processes**) option. For further details, see the **ss(8)** man page.

- You can use the new **bridge fdb flush** command to remove specific forwarding database (fdb) entries which match a supplied option. For further details, see the **bridge(8)** man page.

Jira:RHEL-424[1]

Security improvement of the default nftables service configuration

This enhancement adds the **do_masquerade** chain to the default **nftables** service configuration in the **/etc/sysconfig/nftables/nat.nft** file. This reduces the risk of a port shadow attack, which is described in CVE-2021-3773. The first rule in the **do_masquerade** chain detects suitable packets and enforces source port randomization to reduce the risk of port shadow attacks.

Bugzilla:2061942

NetworkManager supports the no-aaaa DNS option

You can now use the **no-aaaa** option to configure DNS settings on managed nodes by suppressing AAAA queries generated by the stub resolver. Previously, there was no option to suppress AAAA queries generated by the stub resolver, including AAAA lookups triggered by NSS-based interfaces such as **getaddrinfo**; only DNS lookups were affected. With this enhancement, you can disable IPv6 resolution by using the **nmcli** utility. After a restart of the **NetworkManager** service, the **no-aaaa** setting gets reflected in the **/etc/resolv.conf** file, with additional control over DNS lookups.

Bugzilla:2144521

The nm-cloud-setup utility now supports IMDSv2 configuration

Users can configure an AWS Red Hat Enterprise Linux EC2 instance with Instance Metadata Service Version 2 (IMDSv2) with the **nm-cloud-setup** utility. To comply with improved security that restricts unauthorized access to EC2 metadata and new features, integration between AWS and Red Hat services is necessary to provide advanced features. This enhancement enables the **nm-cloud-setup** utility to fetch and save the IMDSv2 tokens, verify an EC2 environment, and retrieve information about available interfaces and IP configuration by using the secured IMDSv2 tokens.

Bugzilla:2151987

The **libnftnl** package rebased to version 1.2.2

The Netlink API to the in-kernel **nf_tables** subsystem (**libnftnl**) package has been rebased. Notable changes and enhancements include:

- Added features:

  - Nesting of the **udata** attribute

  - Resetting TCP options with the **exthdr** expression

  - The **sdif** and **sdifname** meta keywords

  - Support for a new attribute **NFTNL_CHAIN_FLAGS** in the **nftnl_chain** struct, to communicate flags between the kernel and user space.

  - Support for the **nftnl_set** struct nftables sets backend to add expressions to sets and set elements.

  - Comments to sets, tables, objects, and chains

  - The **nftnl_table** struct now has an **NFTNL_TABLE_OWNER** attribute. Set this attribute to enable the kernel to communicate the owner to the user space.

  - Readiness for incremental updates to flowtable device

  - The **typeof** keyword related **nftnl_set udata** definitions

  - The **chain** ID attribute

  - The function to remove expressions from a rule

  - A new **last** expression

- Improved bitwise expressions:

  - Newly added **op** and **data** attributes

  - Left and right shifts

  - Aligned with debug output of other expressions

- Improved socket expressions:

  - Added the **wildcard** attribute

  - Support for cgroups v2

- Improved debug output:

  - Included the **key_end** data register in set elements

  - Dropped unused registers from **masq** and nat expressions

  - Applied fix for verdict map elements

  - Removed leftovers from dropped XML formatting

  - Support for payload offset of inner header

Bugzilla:2211096

4.5. KERNEL

Kernel version in RHEL 8.9

Red Hat Enterprise Linux 8.9 is distributed with the kernel version 4.18.0-513.5.1.

Bugzilla:2232558

The RHEL kernel now supports AutoIBRS

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

Bugzilla:1989283[1]

The Intel® QAT kernel driver rebased to upstream version 6.2

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

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

- Intel Quick Assist Technology 401xx devices

- Intel Quick Assist Technology 402xx devices

Bugzilla:2144529[1]

makedumpfile rebased to version 1.7.2

The **makedumpfile** tool, which makes the crash dump file small by compressing pages or excluding memory pages that are not required, has been rebased to version 1.7.2. The rebase includes many bug fixes and enhancements.

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

Bugzilla:2173791

4.6. FILE SYSTEMS AND STORAGE

Support for specifying a UUID when creating a GFS2 file system

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

Bugzilla:2180782

fuse3 now allows invalidating a directory entry without triggering umount

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

Bugzilla:2171095[1]

4.7. HIGH AVAILABILITY AND CLUSTERS

Pacemaker's scheduler now tries to satisfy all mandatory colocation constraints before trying to satisfy optional colocation constraints

Previously, colocation constraints were considered one by one regardless of whether they were mandatory or optional. This meant that certain resources could be unable to run even though a node assignment was possible. Pacemaker's scheduler now tries to satisfy all mandatory colocation constraints, including the implicit constraints between group members, before trying to satisfy optional colocation constraints. As a result, resources with a mix of optional and mandatory colocation constraints are now more likely to be able to run.

Bugzilla:1876173

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

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

Bugzilla:2040110

The Filesystem resource agent now supports the EFS file system type

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

Bugzilla:2049319

The alert_snmp.sh.sample alert agent now supports SNMPv3

The **alert_snmp.sh.sample** alert agent, which is the sample alert agent provided with Pacemaker, now supports the SNMPv3 protocol as well as SNMPv2. With this update, you can copy the **alert_snmp.sh.sample** agent without modification to use SNMPv3 with Pacemaker alerts.

Bugzilla:2160206

New enabled alert meta option to disable a Pacemaker alert

Pacemaker alerts and alert recipients now support an **enabled** meta option.

- Setting the **enabled** meta option to **false** for an alert disables the alert.

- Setting the **enabled** meta option to **true** for an alert and **false** for a particular recipient disables the alert for that recipient.

The default value for the **enabled** meta option is **true**. You can use this option to temporarily disable an alert for any reason, such as planned maintenance.

Bugzilla:2078611

Pacemaker Remote nodes now preserve transient node attributes after a brief connection outage

Previously, when a Pacemaker Remote connection was lost, Pacemaker would always purge its transient node attributes. This was unnecessary if the connection was quickly recoverable and the remote daemon had not restarted in the meantime. Pacemaker Remote nodes now preserve transient node attributes after a brief, recoverable connection outage.

Bugzilla:2030869

Enhancements to the **pcs property** command

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

- The **pcs property config --output-format=** option

  - Specify **--output-format=cmd** to display the **pcs property set** command created from the current cluster properties configuration. You can use this command to re-create configured cluster properties on a different system.

  - Specify **--output-format=json** to display the configured cluster properties in JSON format.

  - Specify **output-format=text** to display the configured cluster properties in plain text format, which is the default value for this option.

- The **pcs property defaults** command, which replaces the deprecated **pcs property --defaults** option

- The **pcs property describe** command, which describes the meaning of cluster properties.

Bugzilla:2166289

4.8. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

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

A new module stream, **nodejs:20**, previously available as a Technology Preview, is fully supported with the release of the RHEA-2023:7249 advisory. The **nodejs:20** module stream now provides **Node.js 20.9**, which is a Long Term Support (LTS) version.

Node.js 20 included in RHEL 8.9 provides numerous new features, bug fixes, security fixes, and performance improvements over **Node.js 18** available since RHEL 8.7.

Notable changes include:

- The **V8** JavaScript engine has been upgraded to version 11.3.

- The **npm** package manager has been upgraded to version 9.8.0.

- **Node.js** introduces a new experimental Permission Model.

- **Node.js** introduces a new experimental Single Executable Application (SEA) feature.

- **Node.js** provides improvements to the Experimental ECMAScript modules (ESM) loader.

- The native test runner, introduced as an experimental **node:test** module in **Node.js 18**, is now considered stable.

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

```
# yum module install nodejs:20
```

If you want to upgrade from the **nodejs:18** stream, see Switching to a later stream .

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

Bugzilla:2186718

A new **filter** argument to the Python **tarfile** extraction functions

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

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

Jira:RHELDOCS-16405[1]

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

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

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

- **perl-IO-Socket-SSL**

- **perl-Mozilla-CA**

- **perl-Net-SSLeay**

Bugzilla:2228409[1]

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

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

The update in RHSA-2024:0256 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]

4.9. COMPILERS AND DEVELOPMENT TOOLS

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

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

Bugzilla:2180462

GCC now supports preserving register arguments

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

Bugzilla:2168205[1]

New GCC Toolset 13

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

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

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

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

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

```
# yum 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, seehttps://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/developing_c_and_cpp_applications_in_rhel_8/additional-toolsets-for-development_developing-applications#gcc-toolset-13_assembly_additional-toolsets-for-development[GCC Toolset 13] and Using GCC Toolset.

Bugzilla:2171898[1], Bugzilla:2171928, Bugzilla:2188490

GCC Toolset 13: GCC rebased to version 13.1.1

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

General improvements

- OpenMP:

  - OpenMP 5.0: Fortran now supports some non-rectangular loop nests. Such support was added for C/C++ in GCC 11.

  - Many OpenMP 5.1 features have been added.

  - Initial support for OpenMP 5.2 features has been added.

- A new debug info compression option value, **-gz=zstd**, is now available.

- The **-Ofast**, **-ffast-math**, and **-funsafe-math-optimizations** options no longer add startup code to alter the floating-point environment when producing a shared object with the **-shared** option.

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

- Link-time optimization improvements have been implemented.

New languages and language-specific improvements

C family:

- A new **-Wxor-used-as-pow** option warns about uses of the exclusive or (**^**) operator where the user might have meant exponentiation.

- Three new function attributes have been added for documenting **int** arguments that are file descriptors:

  - *attribute*((fd_arg(N)))

- *attribute*((fd_arg_read(N)))

- *attribute*((fd_arg_write(N)))

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

- A new statement attribute, *attribute*((assume(EXPR)));, has been added for C++23 portable assumptions. The attribute is supported also in C or earlier C++.

- GCC can now control when to treat the trailing array of a structure as a flexible array member for the purpose of accessing the elements of such an array. By default, all trailing arrays in aggregates are treated as flexible array members. Use the new command-line option **-fstrict-flex-arrays** to control what array members are treated as flexible arrays.

C:

- Several C23 features have been implemented:

  - Introduced the **nullptr** constant.

  - Enumerations enhanced to specify underlying types.

  - Requirements for variadic parameter lists have been relaxed.

  - Introduced the **auto** feature to enable type inference for object definitions.

  - Introduced the **constexpr** specifier for object definitions.

  - Introduced storage-class specifiers for compound literals.

  - Introduced the **typeof** object (previously supported as an extension) and the **typeof_unqual** object.

  - Added new keywords: **alignas**, **alignof**, **bool**, **false**, **static_assert**, **thread_local**, and **true**.

  - Added the **[[noreturn]]** attribute to specify that a function does not return execution to its caller.

  - Added support for empty initializer braces.

  - Added support for *STDC_VERSION_*_H* header version macros.

  - Removed the **ATOMIC_VAR_INIT** macro.

  - Added the **unreachable** macro for the **<stddef.h>** header.

  - Removed trigraphs.

  - Removed unprototyped functions.

  - Added **printf** and **scanf** format checking through the **-Wformat** option for the **%wN** and **%wfN** format length modifiers.

  - Added support for identifier syntax of Unicode Standard Annex (UAX) 31.

  - Existing features adopted in C23 have been adjusted to follow C23 requirements and are not diagnosed using the **-std=c2x -Wpedantic** option.

- A new **-Wenum-int-mismatch** option warns about mismatches between an enumerated type and an integer type.

C++:

- Implemented excess precision support through the **-fexcess-precision** option. It is enabled by default in strict standard modes like **-std=c++17**, where it defaults to **-fexcess-precision=standard**. In GNU standard modes like **-std=gnu++20**, it defaults to **-fexcess-precision=fast**, which restores previous behavior.
  The **-fexcess-precision** option affects the following architectures:

  - Intel 32- and 64-bit using x87 math, in some cases on Motorola 68000, where **float** and **double** expressions are evaluated in **long double** precision.

  - 64-bit IBM Z systems where **float** expressions are evaluated in **double** precision.

  - Several architectures that support the **std::float16_t** or **std::bfloat16_t** types, where these types are evaluated in **float** precision.

- Improved experimental support for C++23, including:

  - Added support for labels at the end of compound statements.

  - Added a type trait to detect reference binding to a temporary.

  - Reintroduced support for volatile compound operations.

  - Added support for the **#warning** directive.

  - Added support for delimited escape sequences.

  - Added support for named universal character escapes.

  - Added a compatibility and portability fix for the **char8_t** type.

  - Added static **operator()** function objects.

  - Simplified implicit moves.

  - Rewriting equality in expressions is now less of a breaking change.

  - Removed non-encodable wide character literals and wide multicharacter literals.

  - Relaxed some **constexpr** function restrictions.

  - Extended floating-point types and standard names.

  - Implemented portable assumptions.

  - Added support for UTF-8 as a portable source file encoding standard.

  - Added support for static **operator[]** subscripts.

- New warnings:

  - **-Wself-move** warns when a value is moved to itself with **std::move**.

- **-Wdangling-reference** warns when a reference is bound to a temporary whose lifetime has ended.

  - The **-Wpessimizing-move** and **-Wredundant-move** warnings have been extended to warn in more contexts.

- The new **-nostdlib++** option enables linking with **g++** without implicitly linking in the C++ standard library.

Changes in the libstdc++ runtime library

- Improved experimental support for C++20, including:

  - Added the **<format>** header and the **std::format** function.

  - Added support in the **<chrono>** header for the **std::chrono::utc_clock** clock, other clocks, time zones, and the **std::format** function.

- Improved experimental support for C++23, including:

  - Additions to the **<ranges>** header: **views::zip**, **views::zip_transform**, **views::adjacent**, **views::adjacent_transform**, **views::pairwise**, **views::slide**, **views::chunk**, **views::chunk_by**, **views::repeat**, **views::chunk_by**, **views::cartesian_product**, **views::as_rvalue**, **views::enumerate**, **views::as_const**.

  - Additions to the **<algorithm>** header: **ranges::contains**, **ranges::contains_subrange**, **ranges::iota**, **ranges::find_last**, **ranges::find_last_if**, **ranges::find_last_if_not**, **ranges::fold_left**, **ranges::fold_left_first**, **ranges::fold_right**, **ranges::fold_right_last**, **ranges::fold_left_with_iter**, **ranges::fold_left_first_with_iter**.

  - Support for monadic operations for the **std::expected** class template.

  - Added **constexpr** modifiers to the **std::bitset**, **std::to_chars** and **std::from_chars** functions.

  - Added library support for extended floating-point types.

- Added support for the **<experimental/scope>** header from version 3 of the Library Fundamentals Technical Specification (TS).

- Added support for the **<experimental/synchronized_value>** header from version 2 of the Concurrency TS.

- Added support for many previously unavailable features in freestanding mode. For example:

  - The **std::tuple** class template is now available for freestanding compilation.

  - The **libstdc++** library adds components to the freestanding subset, such as **std::array** and **std::string_view**.

  - The **libstdc++** library now respects the **-ffreestanding** compiler option, so it is no longer necessary to build a separate freestanding installation of the **libstdc++** library. Compiling with **-ffreestanding** will restrict the available features to the freestanding subset, even if the **libstdc++** library was built as a full, hosted implementation.

New targets and target-specific Improvements

The 64-bit ARM architecture:

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

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

- For both C and C++, the **__bf16** type is supported on systems with Streaming SIMD Extensions 2 and above enabled.

- The real **__bf16** type is now used for **AVX512BF16** instruction intrinsics. Previously, **__bfloat16**, a typedef of short, was used. Adjust your **AVX512BF16** related source code when upgrading GCC 12 to GCC 13.

- Added new Instruction Set Architecture (ISA) extensions to support the following Intel instructions:

    - **AVX-IFMA** whose instruction intrinsics are available through the **-mavxifma** compiler switch.

    - **AVX-VNNI-INT8** whose instruction intrinsics are available through the **-mavxvnniint8** compiler switch.

    - **AVX-NE-CONVERT** whose instruction intrinsics are available through the **-mavxneconvert** compiler switch.

    - **CMPccXADD** whose instruction intrinsics are available through the **-mcmpccxadd** compiler switch.

    - **AMX-FP16** whose instruction intrinsics are available through the **-mamx-fp16** compiler switch.

    - **PREFETCHI** whose instruction intrinsics are available through the **-mprefetchi** compiler switch.

    - **RAO-INT** whose instruction intrinsics are available through the **-mraoint** compiler switch.

    - **AMX-COMPLEX** whose instruction intrinsics are available through the **-mamx-complex** compiler switch.

- GCC now supports AMD CPUs based on the **znver4** core through the **-march=znver4** compiler switch. The switch makes GCC consider using 512-bit vectors when auto-vectorizing.

Improvements to the static analyzer

- The static analyzer has gained 20 new warnings:

    - **-Wanalyzer-allocation-size**

    - **-Wanalyzer-deref-before-check**

    - **-Wanalyzer-exposure-through-uninit-copy**

    - **-Wanalyzer-imprecise-fp-arithmetic**

    - **-Wanalyzer-infinite-recursion**

    - **-Wanalyzer-jump-through-null**

    - **-Wanalyzer-out-of-bounds**

- **-Wanalyzer-putenv-of-auto-var**

- **-Wanalyzer-tainted-assertion**

- Seven new warnings relating to misuse of file descriptors:

  - **-Wanalyzer-fd-access-mode-mismatch**

  - **-Wanalyzer-fd-double-close**

  - **-Wanalyzer-fd-leak**

  - **-Wanalyzer-fd-phase-mismatch** (for example, calling **accept** on a socket before calling **listen** on it)

  - **-Wanalyzer-fd-type-mismatch** (for example, using a stream socket operation on a datagram socket)

  - **-Wanalyzer-fd-use-after-close**

  - **-Wanalyzer-fd-use-without-check**

    - Also implemented special-casing handling of the behavior of the **open**, **close**, **creat**, **dup**, **dup2**, **dup3**, **pipe**, **pipe2**, **read**, and **write** functions.

- Four new warnings for misuses of the **<stdarg.h>** header:

  - **-Wanalyzer-va-list-leak** warns about missing a **va_end** macro after a **va_start** or **va_copy** macro.

  - **-Wanalyzer-va-list-use-after-va-end** warns about a **va_arg** or **va_copy** macro used on a **va_list** object type that has had the **va_end** macro called on it.

  - **-Wanalyzer-va-arg-type-mismatch** type-checks **va_arg** macro usage in interprocedural execution paths against the types of the parameters that were actually passed to the variadic call.

  - **-Wanalyzer-va-list-exhausted** warns if a **va_arg** macro is used too many times on a **va_list** object type in interprocedural execution paths.

- Numerous other improvements.

Backwards incompatible changes

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

Bugzilla:2172091[1]

GCC Toolset 13: annobin rebased to version 12.20

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

- Added support for moving **annobin** notes into a separate debug info file. This results in reduced executable binary size.

- Added support for a new smaller note format reduces the size of the separate debuginfo files and the time taken to create these files.

Bugzilla:2171923[1]

GCC Toolset 13: GDB rebased to version 12.1

GCC Toolset 13 provides GDB version 12.1.

Notable bug fixes and enhancements include:

- GDB now styles source code and disassembler by default. If styling interferes with automation or scripting of GDB, you can disable it by using the **maint set gnu-source-highlight enabled off** and **maint set style disassembler enabled off** commands.

- GDB now displays backtraces whenever it encounters an internal error. If this affects scripts or automation, you can use the **maint set backtrace-on-fatal-signal off** command to disable this feature.

C/C++ improvements:

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

Terminal user interface (TUI):

- **tui layout**
  tui focus

  tui refresh

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

- **tui window width**
  winwidth

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

- **info win**
  This command now includes information about the width of the TUI windows in its output.

Machine Interface (MI) changes:

- The default version of the MI interpreter is now 4 (–i=mi4).

- The **-add-inferior** command with no flag now inherits the connection of the current inferior. This restores the behavior of GDB prior to version 10.

- The **-add-inferior** command now accepts a **--no-connection** flag that causes the new inferior to start without a connection.

- The **script** field in breakpoint output (which is syntactically incorrect in MI 3 and earlier) has become a list in MI 4. This affects the following commands and events:

  - **-break-insert**

  - **-break-info**

  - **=breakpoint-created**

  - **=breakpoint-modified**
    Use the **-fix-breakpoint-script-output** command to enable the new behavior with earlier MI versions.

New commands:

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

  maint set internal-warning backtrace [on|off]

  maint show internal-warning backtrace

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

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

- **maint set gnu-source-highlight enabled [on|off]**
  maint show gnu-source-highlight enabled
  Enables or disables the GNU Source Highlight library for adding styling to source code. When disabled, the library is not used even if it is available. When the GNU Source Highlight library is not used the Python Pygments library is used instead.

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

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

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

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

Changed commands:

- **set logging [on|off]**
  Deprecated and replaced by the **set logging enabled [on|off]** command.

- **print**

Printing of floating-point values with base-modifying formats like **/x** has been changed to display the underlying bytes of the value in the desired base.

- **clone-inferior**
  The **clone-inferior** command now ensures that the **TTY**, **CMD**, and **ARGs** settings are copied from the original inferior to the new one. All modifications to the environment variables done using the **set environment** or **unset environment** commands are also copied to the new inferior.

Python API:

- The new **gdb.add_history()** function takes a **gdb.Value** object and adds the value it represents to GDB's history list. The function returns an integer, which is the index of the new item in the history list.

- The new **gdb.history_count()** function returns the number of values in GDB's value history.

- The new **gdb.events.gdb_exiting** event is called with a **gdb.GdbExitingEvent** object that has the read-only attribute **exit_code** containing the value of the GDB exit code. This event is triggered prior to GDB's exit before GDB starts to clean up its internal state.

- The new **gdb.architecture_names()** function returns a list containing all of the possible **Architecture.name()** values. Each entry is a string.

- The new **gdb.Architecture.integer_type()** function returns an integer type given a size and a signed-ness.

- The new **gdb.TargetConnection** object type represents a connection (as displayed by the **info connections** command). A sub-class, **gdb.RemoteTargetConnection**, represents **remote** and **extended-remote** connections.

- The **gdb.Inferior** type now has a **connection** property that is an instance of the **gdb.TargetConnection** object, the connection used by this inferior. This can be **None** if the inferior has no connection.

- The new **gdb.events.connection_removed** event registry emits a **gdb.ConnectionEvent** event when a connection is removed from GDB. This event has a **connection** property, a **gdb.TargetConnection** object for the connection being removed.

- The new **gdb.connections()** function returns a list of all currently active connections.

- The new **gdb.RemoteTargetConnection.send_packet(PACKET)** method is equivalent to the existing **maint packet** CLI command. You can use it to send a specified packet to the remote target.

- The new **gdb.host_charset()** function returns the name of the current host character set as a string.

- The new **gdb.set_parameter(*NAME*, *VALUE*)** function sets the GDB parameter *NAME* to *VALUE*.

- The new **gdb.with_parameter(*NAME*, *VALUE*)** function returns a context manager that temporarily sets the GDB parameter *NAME* to *VALUE* and then resets it when the context is exited.

- The **gdb.Value.format_string** method now takes a **styling** argument, which is a boolean. When **true**, the returned string can include escape sequences to apply styling. The styling is present

only if styling is turned on in GDB (see **help set styling**). When **false**, which is the default if the **styling** argument is not given, no styling is applied to the returned string.

- The new read-only attribute **gdb.InferiorThread.details** is either a string containing additional target-specific thread-state information, or **None** if there is no such additional information.

- The new read-only attribute **gdb.Type.is_scalar** is **True** for scalar types, and **False** for all other types.

- The new read-only attribute **gdb.Type.is_signed** should only be read when **Type.is_scalar** is **True**, and will be **True** for signed types and **False** for all other types. Attempting to read this attribute for non-scalar types will raise a **ValueError**.

- You can now add GDB and MI commands implemented in Python.

For more information see the upstream release notes:

What has changed in GDB?

Bugzilla:2172095[1]

GCC Toolset 13: **bintuils** rebased to version 2.40

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

Linkers:

- The new **-w** (**--no-warnings**) command-line option for the linker suppresses the generation of any warning or error messages. This is useful in case you need to create a known non-working binary.

- The ELF linker now generates a warning message if:

  - The stack is made executable

  - It creates a memory resident segment with all three of the **Read**, **Write** and **eXecute** permissions set

  - It creates a thread local data segment with the **eXecute** permission set.
    You can disable these warnings by using the **--no-warn-exec-stack** or **--no-warn-rwx-segments** options.

- The linker can now insert arbitrary JSON-format metadata into binaries that it creates.

Other tools:

- A new the **objdump** tool's **--private** option to display fields in the file header and section headers for Portable Executable (PE) format files.

- A new **--strip-section-headers** command-line option for the **objcopy** and **strip** utilities to remove the ELF section header from ELF files.

- A new **--show-all-symbols** command-line option for the **objdump** utility to display all symbols that match a given address when disassembling, as opposed to the default function of displaying only the first symbol that matches an address.

- A new **-W** (**--no-weak**) option to the **nm** utility to make it ignore weak symbols.

- The **objdump** utility now supports syntax highlighting of disassembler output for some architectures. Use the **--disassembler-color=*MODE*** command-line option, with *MODE* being one of the following:

  - **off**

  - **color** – This option is supported by all terminal emulators.

  - **extended-color** – This option uses 8-bit colors not supported by all terminal emulators.

Bugzilla:2171924[1]

GCC Toolset 13: **annobin** rebased to version 12.20

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

- Added support for moving **annobin** notes into a separate debug info file. This results in reduced executable binary size.

- Added support for a new smaller note format, which reduces the size of the separate debuginfo files and the time taken to create these files.

Bugzilla:2171921[1]

Valgrind rebased to version 3.21.0

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

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

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

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

- The **vgdb** utility now supports the extended remote protocol when invoked with the **--multi** option. The GDB **run** command is supported in this mode and, as a result, you can run GDB and Valgrind from a single terminal.

- You can use the **--realloc-zero-bytes-frees=[yes|no]** option to change the behavior of the **realloc()** function with a size of zero for tools that intercept the **malloc()** call.

- The **memcheck** tool now performs checks for the use of the **realloc()** function with a size of zero. Use the new **--show-realloc-size-zero=[yes|no]** switch to disable this feature.

- You can use the new **--history-backtrace-size=*value*** option for the **helgrind** tool to configure the number of entries to record in the stack traces of earlier accesses.

- The **--cache-sim=[yes|no] cachegrind** option now defaults to **no** and, as a result, only instruction cache read events are gathered by default.

- The source code for the **cg_annotate**, **cg_diff**, and **cg_merge cachegrind** utilities has been

rewritten and, as a result, the utilities have more flexible command line option handling. For example, they now support the **--show-percs** and **--no-show-percs** options as well as the existing **--show-percs=yes** and **--show-percs=no** options.

- The **cg_annotate cachegrind** utility now supports diffing (using the **--diff**, **--mod-filename**, and **--mod-funcname** options) and merging (by passing multiple data files). In addition, **cg_annotate** now provides more information at the file and function level.

- A new user-request for the **DHAT** tool allows you to override the 1024 byte limit on access count histograms for blocks of memory.

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

- 64-bit ARM:

    - v8.2 scalar and vector Floating-point Absolute Difference (FABD), Floating-point Absolute Compare Greater than or Equal (FACGE), Floating-point Absolute Compare Greater Than (FACGT), and Floating-point Add (FADD) instructions.

    - v8.2 Floating-point (FP) compare and conditional compare instructions.

    - Zero variants of v8.2 Floating-point (FP) compare instructions.

- 64-bit IBM Z:

    - Support for the **miscellaneous-instruction-extensions facility 3** and the **vector-enhancements facility 2**. This enables programs compiled with the **-march=arch13** or **-march=z15** options to be executed under Valgrind.

- IBM Power:

    - ISA 3.1 support is now complete.

    - ISA 3.0 now supports the deliver a random number (darn) instruction.

    - ISA 3.0 now supports the System Call Vectored (scv) instruction.

    - ISA 3.0 now supports the copy, paste, and cpabort instructions.

Bugzilla:2124345

systemtap rebased to version 4.9

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

- A new Language-Server-Protocol (LSP) backend for easier interactive drafting of **systemtap** scripts on LSP-capable editors.

- Access to a Python/Jupyter interactive notebook frontend.

- Improved handling of DWARF 5 bitfields.

Bugzilla:2186932

elfutils rebased to version 0.189

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

libelf

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

libdwfl

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

eu-elfcompress

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

Bugzilla:2182060

libpfm rebased to version 4.13

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

- AMD Zen 4

- ARM Neoverse N1

- ARM Neoverse N2

- ARM Neoverse V1

- ARM Neoverse V2

- 4th Generation Intel® Xeon® Scalable Processors

- IBM z16

Bugzilla:2185653, Bugzilla:2111987, Bugzilla:2111966, Bugzilla:2111973, Bugzilla:2109907, Bugzilla:2111981, Bugzilla:2047725

papi supports new processor microarchitectures

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

- ARM Neoverse N1

- ARM Neoverse N2

- ARM Neoverse V1

- ARM Neoverse V2

Bugzilla:2111982[1], Bugzilla:2111988

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

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

Bugzilla:2161146[1]

LLVM Toolset rebased to version 16.0.6

LLVM Toolset has been updated to version 16.0.6.

Notable enhancements include:

- Improvements to optimization

- Support for new CPU extensions

- Improved support for new C++ versions.

Notable backwards incompatible changes include:

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

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

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

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

Bugzilla:2178806

Rust Toolset rebased to version 1.71.1

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

- A new implementation of multiple producer, single consumer (mpsc) channels to improve performance

- A new Cargo **sparse** index protocol for more efficient use of the **crates.io** registry

- New **OnceCell** and **OnceLock** types for one-time value initialization

- A new **C-unwind** ABI string to enable usage of forced unwinding across Foreign Function Interface (FFI) boundaries

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

- Announcing Rust 1.67.0

- Announcing Rust 1.68.0

- Announcing Rust 1.69.0

- Announcing Rust 1.70.0

- Announcing Rust 1.71.0

Bugzilla:2191740

The Rust profiler_builtins runtime component is now available

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

-C instrument-coverage

Enables coverage profiling

-C profile-generate

Enables profile-guided optimization

Bugzilla:2213875[1]

Go Toolset rebased to version 1.20.10

Go Toolset has been updated to version 1.20.10.

Notable enhancements include:

- New functions added in the **unsafe** package to handle slices and strings without depending on the internal representation.

- Comparable types can now satisfy comparable constraints.

- A new **crypto/ecdh** package.

- The **go build** and **go test** commands no longer accept the **-i** flag.

- The **go generate** and **go test** commands now accept the **-skip** *pattern* option.

- The **go build**, **go install**, and other build-related commands now support the **-pgo** and **-cover** flags.

- The **go** command now disables **cgo** by default on systems without a C toolchain.

- The **go version -m** command now supports reading more Go binaries types.

- The **go** command now disables **cgo** by default on systems without a C toolchain.

- Added support for collecting code coverage profiles from applications and integration tests instead of collecting them only from unit tests.

Bugzilla:2185260[1]

grafana rebased to version 9.2.10

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

- The time series panel is now the default visualization option, replacing the graph panel.

- Grafana provides a new Prometheus and Loki query builder.

- Grafana now includes multiple UI/UX and performance improvements.

- The license has changed from Apache 2.0 to GNU Affero General Public License (AGPL).

- The heatmap panel is now used throughout Grafana.

- Geomaps can now measure both distance and area.

- The Alertmanager is now based on **Prometheus Alertmanager** version 0.24.

- Grafana Alerting rules now return an **Error** state by default on execution error or timeout.

- Expressions can now be used on public dashboards.

- The join transformation now supports inner joins.

- Public dashboards now allow sharing Grafana dashboards.

- A new Prometheus streaming parser is now available as an opt-in feature.

For more information, see the upstream release notes:

- What's new in Grafana v8.0

- What's new in Grafana v9.0

- What's new in Grafana v9.1

- What's new in Grafana v9.2

Bugzilla:2193250

grafana-pcp rebased to version 5.1.1

The **grafana-pcp** package, which provides the Performance Co-Pilot Grafana Plugin, has been updated to version 5.1.1. Notable changes include:

- Query editor: Added buttons to disable rate conversation and time utilization conversation

- Redis datasource:

  - Removed the deprecated **label_values(metric, label)** function

  - Fixed the network error for metrics with many series (requires Performance Co-Pilot version 6 and later)

- Set the **pmproxy** API timeout to 1 minute

Bugzilla:2193270

.NET 8.0 is available

Red Hat Enterprise Linux 8.9 is distributed with .NET version 8.0. Notable improvements include:

- Added support for the C#12 and F#8 language versions.

- Added support for building container images using the .NET Software Development Kit directly.

- Many performance improvements to the garbage collector (GC), Just-In-Time (JIT) compiler, and the base libraries.

Jira:RHELPLAN-164398[1]

4.10. IDENTITY MANAGEMENT

samba rebased to version 4.18.4

The **samba** packages have been upgraded to upstream version 4.18.4, which provides bug fixes and enhancements over the previous version. The most notable changes:

- Security improvements in previous releases impacted the performance of the Server Message Block (SMB) server for high metadata workloads. This update improves the performance in this scenario.

- The new **wbinfo --change-secret-at=<domain_controller>** command enforces the change of the trust account password on the specified domain controller.

- By default, Samba stores access control lists (ACLs) in the **security.NTACL** extended attribute of files. You can now customize the attribute name with the **acl_xattr:<security_acl_name>** setting in the **/etc/samba/smb.conf** file. Note that a custom extended attribute name is not a protected location as **security.NTACL**. Consequently, users with local access to the server can be able to modify the custom attribute's content and compromise the ACL.

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

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

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

Bugzilla:2190417

ipa rebased to version 4.9.12

The **ipa** package has been upgraded to version 4.9.12. For more information, see the upstream FreeIPA release notes.

Bugzilla:2196425

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

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

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

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

Jira:RHELDOCS-16474[1]

ansible-freeipa ipaserver role now supports Random Serial Numbers

With this update, you can use the **ipaserver_random_serial_numbers=true** option with the **ansible-freeipa ipaserver** role. This way, you can generate fully random serial numbers for certificates and requests in PKI when installing an Identity Management (IdM) server using Ansible. With RSNv3, you can avoid range management in large IdM installations and prevent common collisions when reinstalling IdM.
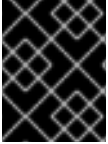
> **IMPORTANT**
>
> RSNv3 is supported only for new IdM installations. If enabled, it is required to use RSNv3 on all PKI services.

Jira:RHELDOCS-16462[1]

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

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

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

- Set **ipaserver_ignore_topology_disconnect** to True.

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

Bugzilla:2127901

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

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

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

Bugzilla:2175766

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

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

Bugzilla:2127906

MIT Kerberos now supports the Extended KDC MS-PAC signature

With this update, MIT Kerberos, which is used by Red Hat, implements support for one of the two types of the Privilege Attribute Certificate (PAC) signatures introduced by Microsoft in response to recent CVEs. Specifically, MIT Kerberos in RHEL 8 supports the Extended KDC signature that was released in KB5020805 and that addresses CVE-2022-37967.

Note that because of ABI stability constraints, MIT Kerberos on RHEL8 cannot support the other PAC signature type, that is Ticket signature as defined in KB4598347.

To troubleshoot problems related to this enhancement, see the following Knowledgebase resources:

- RHEL-8.9 IdM update, web UI and CLI 401 Unauthorized with KDC S4U2PROXY_EVIDENCE_TKT_WITHOUT_PAC - user and group objects need SIDs

- find_sid_for_ldap_entry - [file ipa_sidgen_cofind_sid_for_ldap_entry - [file ipa_sidgen_common.c, line 521]: Cannot convert Posix ID [120000023l] into an unused SID]

- When upgrading to RHEL9, IDM users are not able to login anymore

- POSIX IDs, SIDs and IDRanges in IPA

See also BZ#2211387 and BZ#2176406.

Bugzilla:2211390

RHEL 8.9 provides **389-ds-base** 1.4.3.37

RHEL 8.9 is distributed with the **389-ds-base** package version 1.4.3.37.

Bugzilla:2188628

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

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

Bugzilla:2166332

4.11. GRAPHICS INFRASTRUCTURES

Intel Arc A-Series graphics is now fully supported

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

With this release, you no longer have to set the **i915.force_probe** kernel option, and full support for these GPUs is enabled by default.

Bugzilla:2041686[1]

4.12. THE WEB CONSOLE

Podman health check action is now available

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

- No action (default): Take no action.

- Restart: Restart the container.

- Stop: Stop the container.

- Force stop: Force stops the container, it does not wait for the container to exit.

Jira:RHELDOCS-16247[1]

Accounts page updates for the web console

This update introduces the following updates to the **Accounts** page:

- It is now possible to add custom user ID and define home directory and shell during the account creation process.

- When creating an account, password validation actively performs a check on every keystroke. Additionally, weak passwords are now shown with a warning.

- Account detail pages now show the home directory and shell for an account.

- It is possible to change shell from the account details page.

Jira:RHELDOCS-16367[1]

4.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The `postgresql` RHEL System Role is now available

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

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

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

Bugzilla:2151371

keylime_server RHEL System Role

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

Bugzilla:2224387

Support for new `ha_cluster` System Role features

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

- Configuration of resource and resource operation defaults, including multiple sets of defaults with rules.

- Loading and blocking of SBD watchdog kernel modules. This makes installed hardware watchdogs available to the cluster.

- Assignment of distinct passwords to the cluster hosts and the quorum device. With that, you can configure a deployment where the same quorum hosts are joined to multiple, separate clusters, and the passwords of the **hacluster** user on these clusters are different.

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

Bugzilla:2190483, Bugzilla:2190478, Bugzilla:2216485

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

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

Bugzilla:2141961

podman RHEL System Role now supports Quadlets, healthchecks, and secrets

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

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

You can use the **podman_secrets** variable to manage secrets. For details, see upstream documentation.

Jira:RHELPLAN-154440[1]

RHEL System Roles now have new volume options for mount point customization

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

Bugzilla:2181661

kdump RHEL System Role updates

The **kdump** RHEL System Role has been updated to a newer version, which brings the following notable enhancements:

- After installing **kexec-tools**, the utility suite no longer generates the **/etc/sysconfig/kdump** file because you do not need to manage this file anymore.

- The role supports the **auto_reset_crashkernel** and **dracut_args** variables.

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

Bugzilla:2211272

The ad_integration RHEL System Role can now rejoin an AD domain

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

Bugzilla:2211723

The rhc System Role now supports setting a proxy server type

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

Bugzilla:2211778

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

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

Bugzilla:2216759

The **certificate** RHEL System Role now allows changing certificate file mode when using certmonger

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

Bugzilla:2218204

New RHEL System Role for managing `systemd` units

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

Bugzilla:2224388

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

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

Bugzilla:2218595

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

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

Bugzilla:2211273

firewall RHEL System Role supports variables related to `ipsets`

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

You can manage **ipsets** with the **firewall** RHEL System Role using the following variables:

- **ipset**

- **ipset_type**

- **ipset_entries**

- **short**

- **description**

- **state: present** or **state: absent**

- **permanent: true**

The following are some notable benefits of this enhancement:

- You can reduce the complexity of the rich rules that define rules for many IP addresses.

- You can add or remove IP addresses from sets as needed without modifying multiple rules.

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

Bugzilla:2140880

Improved performance of the **selinux** System Role with **restorecon -T 0**

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

Bugzilla:2192343

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

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

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

Bugzilla:2222809

The **podman** RHEL System Role now uses **getsubids** to get subuids and subgids

The **podman** RHEL System role now uses the **getsubids** command to get the subuid and subgid ranges for a user and group, respectively. The **podman** RHEL System role also uses this command to verify users and groups to work with identity management.

Jira:RHEL-866[1]

The **podman_kube_specs** variable now supports **pull_image** and **continue_if_pull_fails** fields

The **podman_kube_specs** variable now supports new fields:

- **pull_image**: ensures the image is pulled before use. The default value is **true**. Use **false** if you have some other mechanism to ensure the images are present on the system and you do not want to pull the images.

- **continue_if_pull_fails**: If pulling image fails, it is not treated as a fatal error, and continues with the role. The default is **false**. Use **true** if you have some other mechanism to ensure the correct images are present on the system.

Jira:RHEL-858[1]

Resetting the firewall RHEL System Role configuration now requires minimal downtime

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

Bugzilla:2224648

4.14. RHEL IN CLOUD ENVIRONMENTS

cloud-init supports NetworkManager keyfiles

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

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

```
# cat /etc/cloud/cloud.cfg

network:
    renderers: ['network-manager', 'eni', 'netplan', 'sysconfig', 'networkd']
```

Bugzilla:2219528[1]

cloud-init now uses VMware datasources by default on ESXi

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

Bugzilla:2230777[1]

4.15. SUPPORTABILITY

sos rebased to version 4.6

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

- **sos** reports now include the contents of both **/boot/grub2/custom.cfg** and **/boot/grub2/user.cfg** files that might contain critical information for troubleshooting boot issues. (BZ#2213951)

- The **sos** plugin for OVN-Kubernetes collects additional logs for the interconnect environment. With this update, **sos** also collects logs from the **ovnkube-controller** container when both **ovnkube-node** and **ovnkube-controller** containers are merged into one.

In addition, notable bug fixes include:

- **sos** now correctly gathers **cgroup** data in the OpenShift Container Platform 4 environment (BZ#2186361).

- While collecting **sos** reports with the **sudo** plugin enabled, **sos** now removes the **bindpw** option properly. (BZ#2143272)

- The **subscription_manager** plugin no longer collects proxy usernames and passwords from the **/var/lib/rhsm/** path. (BZ#2177282)

- The **virsh** plugin no longer collects the SPICE remote-display passwords in virt-manager logs, which prevents **sos** from disclosing passwords in its reports. (BZ#2184062)

- **sos** now masks usernames and passwords previously displayed in the **/var/lib/iscsi/nodes/<IQN>/<PortalIP>/default** file.
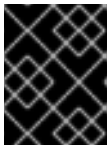
> **IMPORTANT**
>
> The generated archive might contain data considered sensitive. Thus, you should always review the content before passing it to any third party.

   (BZ#2187859)

- **sos** completes the tailed log collection even when the size of the log file is exceeded and when a plugin times out. (BZ#2203141)

- When entering the **sos collect** command on a Pacemaker cluster node, **sos** collects an sos report from the same cluster node. (BZ#2186460)

- When collecting data from a host in the OpenShift Container Platform 4 environment, **sos** now uses the **sysroot** path, which ensures that only the correct data are assembled. (BZ#2075720)

- The **sos report --clean** command obfuscates all MAC addresses as intended. (BZ#2207562)

- Disabling the **hpssm** plugin no longer raises exceptions. (BZ#2216608)

- The **sos clean** command follows permissions of sanitized files. (BZ#2218279)

For details on each release of **sos**, see upstream release notes.

Jira:RHELPLAN-156196[1]

4.16. CONTAINERS

Podman supports pulling and pushing images compressed with zstd

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

Jira:RHELPLAN-154313[1]

Quadlet in Podman is now available

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

For more details, see the Quadlet upstream documentation and the Make systemd better for Podman with Quadlet article.

Jira:RHELPLAN-154431[1]

The Container Tools packages have been updated

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

Notable changes in Podman v4.6 include:

- The **podman kube play** command now supports the **--configmap=_<path>_** option to provide Kubernetes YAML file with environment variables used within the containers of the pod.

- The **podman kube play** command now supports multiple Kubernetes YAML files for the **--configmap** option.

- The **podman kube play** command now supports containerPort names and port numbers within liveness probes.

- The **podman kube play** command now adds the ctrName as an alias to the pod network.

- The **podman kube play** and **podman kube generate** commands now support SELinux filetype labels and ulimit annotations.

- A new command, **podman secret exists**, has been added, which verifies if a secret with the given name exists.

- The **podman create**, **podman run**, **podman pod create**, and **podman pod clone** commands now support a new option, **--shm-size-systemd**, which allows limiting tmpfs sizes for systemd-specific mounts.

- The **podman create** and podman run commands now support a new option, **--security-opt label=nested**, which allows SELinux labeling within a confined container.

- Podman now supports auto updates for containers running inside a pod.

- Podman can now use an SQLite database as a backend for increased stability. The default remains the BoltDB database. You can select the database by setting the **database_backend** field in the **containers.conf** file.

- Podman now supports Quadlets to automatically generate a **systemd** service file from the container description. The description focuses on the relevant container details and hides the technical complexity of running containers under **systemd**.

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

Jira:RHELPLAN-154443[1]

Podman now supports a Podmansh login shell

Beginning with Podman v4.6, you can use the **Podmansh** login shell to manage user access and control. To switch to CGroups v2, add **systemd.unified_cgroup_hierarchy=1** to the kernel command line. Configure the settings for a user to use the **/usr/bin/podmansh** command as a login shell instead of a standard shell command, for example, **/usr/bin/bash**. When a user logs into a system setup, the **podmansh** command runs the user's session in a Podman container named **podmansh**. Containers into which users log in are defined using the Quadlet files, which are created in the **/etc/containers/systemd/users/** directory. In these files, set the **ContainerName** field in the **[Container]** section to **podmansh**. Systemd automatically starts **podmansh** when the user session starts and continues running until all user sessions exit.

For more information, see Podman v4.6.0 Introduces Podmansh: A Revolutionary Login Shell .

Jira:RHELPLAN-163002[1]

Clients for sigstore signatures with Fulcio and Rekor are now available

With Fulcio and Rekor servers, you can now create signatures by using short-term certificates based on an OpenID Connect (OIDC) server authentication, instead of manually managing a private key. Clients for sigstore signatures with Fulcio and Rekor, previously available as a Technology Preview, are now fully supported. This added functionality is the client side support only, and does not include either the Fulcio or Rekor servers.

Add the **fulcio** section in the **policy.json** file. To sign container images, use the **podman push --sign-by-sigstore=file.yml** or **skopeo copy --sign-by-sigstore=_file.yml_** commands, where _file.yml_ is the sigstore signing parameter file.

To verify signatures, add the **fulcio** section and the **rekorPublicKeyPath** or **rekorPublicKeyData** fields in the **policy.json** file. For more information, see **containers-policy.json** man page.

Jira:RHELPLAN-160659[1]

CHAPTER 5. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS

This chapter provides system administrators with a summary of significant changes in the kernel shipped with Red Hat Enterprise Linux 8.9. 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

gather_data_sampling=[X86,INTEL]

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

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

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

- **off**: Disable GDS mitigation.

rdrand=[X86]

With this kernel parameter, you can hide the advertisement of RDRAND support. This affects certain AMD processors because of buggy BIOS support, specifically around the suspend or resume path.

- **force**: Override the decision by the kernel to hide the advertisement of RDRAND support.

Updated kernel parameters

intel_pstate=[X86]

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

- **disable** - Do not enable **intel_pstate** as the default scaling driver for the supported processors.

- **[NEW] active** - Use **intel_pstate** driver to bypass the scaling governors layer of **cpufreq** and provides it own algorithms for p-state selection. There are two P-state selection algorithms provided by **intel_pstate** in the active mode: powersave and performance. The way they both operate depends on whether or not the hardware managed P-states (HWP) feature has been enabled in the processor and possibly on the processor model.

- **passive** - Use **intel_pstate** as a scaling driver, but configure it to work with generic **cpufreq** governors (instead of enabling its internal governor). This mode cannot be used along with the hardware-managed P-states (HWP) feature.

- **force** - Enable **intel_pstate** on systems that prohibit it by default in favor of **acpi-cpufreq**. Forcing the **intel_pstate** driver instead of **acpi-cpufreq** might disable platform features, such as thermal controls and power capping, that rely on ACPI P-States information being

indicated to OSPM and therefore should be used with caution. This option does not work with processors that are not supported by the **intel_pstate** driver or on platforms that use **pcc-cpufreq** instead of **acpi-cpufreq**.

- **no_hwp** - Do not enable hardware P state control (HWP) if available.

- **hwp_only** - Only load **intel_pstate** on systems that support hardware P state control (HWP) if available.

- **support_acpi_ppc** - Enforce **ACPI _PPC** performance limits. If the Fixed ACPI Description Table specifies preferred power management profile as "Enterprise Server" or "Performance Server", then this feature is turned on by default.

- **per_cpu_perf_limits** - Allow per-logical-CPU P-State performance control limits using the **cpufreq sysfs** interface.

rdt=[HW,X86,RDT]

With this kernel parameter, you can turn on or off individual RDT features. The list includes: **cmt**, **mbmtotal**, **mbmlocal**, **l3cat**, **l3cdp**, **l2cat**, **l2cdp**, **mba**, [NEW] **smba**, [NEW] **bmec**.
For example, to turn on **cmt** and turn off **mba** use:

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

tsc=[x86]

With this kernel parameter, you can disable clocksource stability checks for TSC. This parameter takes the format of: **<string>**.

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

- **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 can add overhead.

- **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.

- **nowatchdog**: disable clocksource watchdog. Used in situations with strict latency requirements (where interruptions from clocksource watchdog are not acceptable).

- **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 sysctl parameters

nmi_wd_lpm_factor=(PPC only)

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

- A value of **0** means no change.

- The default value is **200** meaning the NMI watchdog is set to 30s (based on **watchdog_thresh** equal to 10).

txrehash

With this kernel parameter, you can control default hash rethink behaviour on socket.

- If set to **1** (default), hash rethink is performed on listening socket.

- If set to **0**, hash rethink is not performed.

CHAPTER 6. DEVICE DRIVERS

6.1. NEW DRIVERS

Network drivers

- Thunderbolt/USB4 network driver (**thunderbolt_net**)

- Broadcom 802.11 wireless LAN fullmac driver (**brcmfmac**) (only in 64-bit ARM architecture)

Graphics drivers and miscellaneous drivers

- Bluetooth support for MediaTek devices ver 0.1 (**btmtk**), only in IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- DRM Buddy Allocator (**drm_buddy**), only in 64-bit IBM Z architecture

- DRM display adapter helper (**drm_display_helper**), only in 64-bit IBM Z architecture

- Microsoft Azure Network Adapter IB driver (**mana_ib**), only in AMD and Intel 64-bit architectures

- The Linux USB Video Class driver (**uvc**), (only in IBM Power Systems, Little Endian and AMD and Intel 64-bit architectures

- Intel Meteor Lake PCH pinctrl/GPIO driver (**pinctrl-meteorlake**), only in AMD and Intel 64-bit architectures

- Intel In Field Scan (IFS) device (**intel_ifs**), only in AMD and Intel 64-bit architectures

- Intel Uncore Frequency Common Module (**intel-uncore-frequency-common**), only in AMD and Intel 64-bit architectures

- Intel Uncore Frequency Limits Driver (**intel-uncore-frequency**), only in AMD and Intel 64-bit architectures

- AMD SoundWire driver (**soundwire-amd**), only in AMD and Intel 64-bit architectures

- DisplayPort Alternate Mode (**typec_displayport**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Virtio-mem driver (**virtio_mem**), only in AMD and Intel 64-bit architectures

6.2. UPDATED DRIVERS

Network driver updates

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

The following drivers have been updated to **4.18.0-513.5.1 kernel version**:

- Intel® 10 Gigabit PCI Express Network Driver (**ixgbe**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® 10 Gigabit Virtual Function Network Driver (**ixgbevf**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® 2.5G Ethernet Linux Driver (**igc**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® Ethernet Adaptive Virtual Function Network Driver (**iavf**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® Ethernet Connection XL710 Network Driver (**i40e**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® Ethernet Switch Host Interface Driver (**fm10k**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® Gigabit Ethernet Network Driver (**igb**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® Gigabit Virtual Function Network Driver (**igbvf**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Intel® PRO/1000 Network Driver (**e1000e**), only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures

- Mellanox 5th generation network adapters (ConnectX series) core driver (**mlx5_core**)

- The Netronome Flow Processor (NFP) driver (**nfp**)

Graphics, storage, and miscellaneous driver updates

- Broadcom MegaRAID SAS Driver (**megaraid_sas**) has been updated to version 07.725.01.00-rc1, (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)

- Driver for Microchip Smart Family Controller version (**smartpqi**) has been updated to version 2.1.22-040 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)

- Emulex LightPulse Fibre Channel SCSI driver (**lpfc**) has been updated to version 0:14.0.0.21 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)

- MPI3 Storage Controller Device Driver (**mpi3mr**) has been updated to version 8.4.1.0.0 (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)

- QLogic Fibre Channel HBA Driver (**qla2xxx**) has been updated to version 10.02.08.200-k (only in 64-bit ARM architecture, IBM Power Systems, Little Endian, and AMD and Intel 64-bit architectures)

CHAPTER 7. AVAILABLE BPF FEATURES

This chapter provides the complete list of **Berkeley Packet Filter** (**BPF**) features available in the kernel of this minor version of Red Hat Enterprise Linux 8. 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	1 (bpf() syscall restricted to privileged users, without recovery)
JIT compiler	1 (enabled)
JIT compiler hardening	1 (enabled for unprivileged users)
JIT compiler kallsyms exports	1 (enabled for root)
Memory limit for JIT for unprivileged users	264241152
CONFIG_BPF	y
CONFIG_BPF_SYSCALL	y
CONFIG_HAVE_EBPF_JIT	y
CONFIG_BPF_JIT	y
CONFIG_BPF_JIT_ALWAYS_ON	y
CONFIG_DEBUG_INFO_BTF	y
CONFIG_DEBUG_INFO_BTF_MODULES	n
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	y
CONFIG_NET	y
CONFIG_XDP_SOCKETS	y
CONFIG_LWTUNNEL_BPF	y
CONFIG_NET_ACT_BPF	m
CONFIG_NET_CLS_BPF	m
CONFIG_NET_CLS_ACT	y
CONFIG_NET_SCH_INGRESS	m
CONFIG_XFRM	y
CONFIG_IP_ROUTE_CLASSID	y
CONFIG_IPV6_SEG6_BPF	n
CONFIG_BPF_LIRC_MODE2	n
CONFIG_BPF_STREAM_PARSER	y
CONFIG_NETFILTER_XT_MATCH_BPF	m
CONFIG_BPFILTER	n
CONFIG_BPFILTER_UMH	n

Option	Value
CONFIG_TEST_BPF	m
CONFIG_HZ	1000
bpf() syscall	available
Large program size limit	available

Table 7.2. Available program types and supported helpers

Program type	Available helpers
socket_filter	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_load_bytes_relative, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
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_override_return, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
sched_cls	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_change_proto, bpf_skb_change_type, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_skb_adjust_room, bpf_skb_get_xfrm_state, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_skb_cgroup_id, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_skb_cgroup_classid, bpf_redirect_neigh, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_redirect_peer, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf
sched_act	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_change_proto, bpf_skb_change_type, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_skb_adjust_room, bpf_skb_get_xfrm_state, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_skb_cgroup_id, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_skb_cgroup_classid, bpf_redirect_neigh, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_redirect_peer, bpf_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf

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_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
xdp	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_redirect, bpf_perf_event_output, bpf_csum_diff, bpf_get_current_task, bpf_get_numa_node_id, bpf_xdp_adjust_head, bpf_redirect_map, bpf_xdp_adjust_meta, bpf_xdp_adjust_tail, bpf_fib_lookup, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_skc_lookup_tcp, bpf_tcp_check_syncookie, bpf_tcp_gen_syncookie, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_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_ktime_get_coarse_ns, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf
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_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_read_branch_records, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
cgroup_skb	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_load_bytes_relative, bpf_skb_cgroup_id, bpf_get_local_storage, bpf_skb_ancestor_cgroup_id, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_fullsock, bpf_tcp_sock, bpf_skb_ecn_set_ce, bpf_get_listener_sock, bpf_skc_lookup_tcp, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_sk_cgroup_id, bpf_sk_ancestor_cgroup_id, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_sock	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sk_storage_get, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lwt_in	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_cgroup_classid, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_lwt_push_encap, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
lwt_out	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_cgroup_classid, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lwt_xmit	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_cgroup_classid, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_get_tunnel_opt, bpf_skb_set_tunnel_opt, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_csum_update, bpf_set_hash_invalid, bpf_get_numa_node_id, bpf_skb_change_head, bpf_lwt_push_encap, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_csum_level, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
sock_ops	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_setsockopt, bpf_sock_map_update, bpf_getsockopt, bpf_sock_ops_cb_flags_set, bpf_sock_hash_update, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_tcp_sock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_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_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
sk_skb	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_skb_store_bytes, bpf_tail_call, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_current_task, bpf_skb_change_tail, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_skb_change_head, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_skb_adjust_room, bpf_sk_redirect_map, bpf_sk_redirect_hash, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_skc_lookup_tcp, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_device	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_uid_gid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
sk_msg	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_msg_redirect_map, bpf_msg_apply_bytes, bpf_msg_cork_bytes, bpf_msg_pull_data, bpf_msg_redirect_hash, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_msg_push_data, bpf_msg_pop_data, bpf_spin_lock, bpf_spin_unlock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_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_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

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_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_sock_addr	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_setsockopt, bpf_getsockopt, bpf_bind, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_sk_lookup_tcp, bpf_sk_lookup_udp, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_skc_lookup_tcp, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lwt_seg6local	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_cgroup_classid, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_csum_diff, bpf_skb_under_cgroup, bpf_get_hash_recalc, bpf_get_current_task, bpf_skb_pull_data, bpf_get_numa_node_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
lirc_mode2	not supported

Program type	Available helpers
sk_reuseport	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_socket_cookie, bpf_skb_load_bytes_relative, bpf_sk_select_reuseport, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
flow_dissector	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_skb_load_bytes, bpf_get_current_task, bpf_get_numa_node_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
cgroup_sysctl	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_uid_gid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_sysctl_get_name, bpf_sysctl_get_current_value, bpf_sysctl_get_new_value, bpf_sysctl_set_new_value, bpf_strtol, bpf_strtoul, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
raw_tracepoint_writable	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_send_signal, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_send_signal_thread, bpf_jiffies64, bpf_get_ns_current_pid_tgid, bpf_get_current_ancestor_cgroup_id, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_get_task_stack, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_get_current_task_btf, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Program type	Available helpers
cgroup_sockopt	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_uid_gid, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_tcp_sock, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_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_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf
tracing	not supported

Program type	Available helpers
struct_ops	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_skb_store_bytes, bpf_l3_csum_replace, bpf_l4_csum_replace, bpf_tail_call, bpf_clone_redirect, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_get_cgroup_classid, bpf_skb_vlan_push, bpf_skb_vlan_pop, bpf_skb_get_tunnel_key, bpf_skb_set_tunnel_key, bpf_perf_event_read, bpf_redirect, bpf_get_route_realm, bpf_perf_event_output, bpf_skb_load_bytes, bpf_get_stackid, 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_current_task_under_cgroup, 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_xdp_adjust_head, bpf_probe_read_str, bpf_get_socket_cookie, bpf_get_socket_uid, bpf_set_hash, bpf_setsockopt, bpf_skb_adjust_room, bpf_redirect_map, bpf_sk_redirect_map, bpf_sock_map_update, bpf_xdp_adjust_meta, bpf_perf_event_read_value, bpf_perf_prog_read_value, bpf_getsockopt, bpf_override_return, bpf_sock_ops_cb_flags_set, bpf_msg_redirect_map, bpf_msg_apply_bytes, bpf_msg_cork_bytes, bpf_msg_pull_data, bpf_bind, bpf_xdp_adjust_tail, bpf_skb_get_xfrm_state, bpf_get_stack, bpf_skb_load_bytes_relative, bpf_fib_lookup, bpf_sock_hash_update, bpf_msg_redirect_hash, bpf_sk_redirect_hash, bpf_lwt_push_encap, bpf_lwt_seg6_store_bytes, bpf_lwt_seg6_adjust_srh, bpf_lwt_seg6_action, bpf_rc_repeat, bpf_rc_keydown, bpf_skb_cgroup_id, bpf_get_current_cgroup_id, bpf_get_local_storage, bpf_sk_select_reuseport, 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_msg_push_data, bpf_msg_pop_data, bpf_rc_pointer_rel, 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_sysctl_get_name, bpf_sysctl_get_current_value, bpf_sysctl_get_new_value, bpf_sysctl_set_new_value, bpf_strtol, bpf_strtoul, bpf_sk_storage_get, bpf_sk_storage_delete, bpf_send_signal, bpf_tcp_gen_syncookie, bpf_skb_output, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_tcp_send_ack, bpf_send_signal_thread, bpf_jiffies64, bpf_read_branch_records, bpf_get_ns_current_pid_tgid, bpf_xdp_output, bpf_get_netns_cookie, bpf_get_current_ancestor_cgroup_id, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_seq_printf, bpf_seq_write, 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_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_get_task_stack, bpf_load_hdr_opt, bpf_store_hdr_opt, bpf_reserve_hdr_opt, bpf_inode_storage_get, bpf_inode_storage_delete, bpf_d_path, bpf_copy_from_user, bpf_snprintf_btf, bpf_seq_printf_btf, bpf_skb_cgroup_classid, bpf_redirect_neigh, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_redirect_peer, bpf_task_storage_get, bpf_task_storage_delete, bpf_get_current_task_btf, bpf_bprm_opts_set, bpf_ktime_get_coarse_ns, bpf_ima_inode_hash, bpf_sock_from_file, bpf_check_mtu, bpf_for_each_map_elem, bpf_snprintf, bpf_sys_bpf, bpf_btf_find_by_name_kind, bpf_sys_close
ext	not supported
lsm	not supported

Program type	Available helpers
sk_lookup	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_perf_event_output, bpf_get_current_task, bpf_get_numa_node_id, bpf_sk_release, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_spin_lock, bpf_spin_unlock, bpf_probe_read_user, bpf_probe_read_kernel, bpf_probe_read_user_str, bpf_probe_read_kernel_str, bpf_jiffies64, bpf_sk_assign, bpf_ktime_get_boot_ns, bpf_ringbuf_output, bpf_ringbuf_reserve, bpf_ringbuf_submit, bpf_ringbuf_discard, bpf_ringbuf_query, bpf_skc_to_tcp6_sock, bpf_skc_to_tcp_sock, bpf_skc_to_tcp_timewait_sock, bpf_skc_to_tcp_request_sock, bpf_skc_to_udp6_sock, bpf_snprintf_btf, bpf_per_cpu_ptr, bpf_this_cpu_ptr, bpf_ktime_get_coarse_ns, bpf_for_each_map_elem, bpf_snprintf

Table 7.3. Available map types

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

Map type	Available
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	no
ringbuf	yes
inode_storage	yes
task_storage	no

CHAPTER 8. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 8.9 that have a significant impact on users.

8.1. INSTALLER AND IMAGE CREATION

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

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

Bugzilla:1886985

8.2. SECURITY

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

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

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

Bugzilla:1753646[1]

The automatic screen lock now works correctly even when a USB smart-card reader is removed

Before RHEL 8.9, the **opensc** packages incorrectly handled removing USB smart-card readers. Consequently, the system remained unlocked even if the GNOME Display Manager (GDM) was configured to lock the screen when a smart card was removed. Furthermore, after reconnecting the USB reader, the screen also did not lock after removing the smart card. In this release, the code for handling removals of USB smart-card readers has been fixed. As a result, the screen is correctly locked even when a smart card or a USB smart-card reader is removed.

Bugzilla:2097048

The SCAP enable_fips_mode rule now checks only fips=1 on 64-bit IBM Z architecture

Previously, the SCAP Security Guide rule **enable_fips_mode** did check the contents of the /**boot/grub2/grubenv** file. Consequently, the 64-bit IBM Z architecture did not use /**boot/grub2/grubenv** file for FIPS mode. With this update, the OVAL rule **enable_fips_mode** now test if argument **fips=1** for Linux kernel is present in /**boot/loader/entries/.*.conf** file on 64-bit IBM Z architecture.

Bugzilla:2129100

SCAP journald rules no longer remediate to invalid configuration

Previously, the SCAP Security Guide rules **journald_compress**, **journald_forward_to_syslog**, and **journald_storage** contained a bug in the remediation script which added extra quotes to the respective options within the /**etc/systemd/journald.conf** configuration file. Consequently, the **journald** service

failed to parse the configuration options and ignored them. Therefore, the configuration options were not effective and OpenSCAP reported false pass results. With this update, the rules and remediations scripts have been fixed to not add the extra quotes. The rule now create a valid configuration for **journald**.

Bugzilla:2169857

Images can now be configured with security profiles

SCAP Security Guide rules that configure mount point options have been reworked, and you can now use them also for hardening images when building an operating system image in image builder. As a result, you can now build images with partition configuration aligned with a specific security profile.

Bugzilla:2130185

Removed strict requirements from SSG rules related to AIDE configuration

Previously, the SCAP Security Guide (SSG) rule **aide_build_database** required the existence of both **/var/lib/aide/aide.db.new.gz** and **/var/lib/aide/aide.db.gz** files to pass. Because the **AIDE** utility does not require the **/var/lib/aide/aide.db.new.gz** file, this update removed the corresponding requirement from the **aide_build_database** rule. As a result, the rule requires only the **/var/lib/aide/aide.db.gz** file to pass.

In addition, the SCAP Security Guide rule **aide_periodic_cron_checking** is now less strict on entries in **/etc/cron.daily** and **/etc/cron.weekly** files. You can now schedule the **aide --check** command with additional wrappers while staying compliant with the rule.

Bugzilla:2175684

SCAP rules related to **pam_faillock** have correct descriptions

Previously, the SCAP Security Guide rules related to the **pam_faillock** contained descriptions that were misaligned with some profile values. Consequently, the descriptions were not correct. With this update, the rules descriptions are now using XCCDF variables.

This change affects the following rules:

- **accounts_passwords_pam_faillock_deny**

- **accounts_passwords_pam_faillock_interval**

- **accounts_passwords_pam_faillock_dir**

- **accounts_passwords_pam_faillock_unlock_time**

Bugzilla:2175882

The **file_permissions_efi_user_cfg** SCAP rule no longer fails when **/boot/efi** is mounted

Previously, the default permissions of UEFI files were not accepted. Therefore, it was not possible to change the permissions with the **chmod** command when the **/boot/efi** partition used a virtual file allocation table (VFAT) file system. Consequently, the **file_permissions_efi_user_cfg** rule failed. This update changes the default permissions from **0600** to **0700**. Because the **0700** permission is also accepted by CIS, the assessment and remediation are now better aligned with CIS profiles.

Bugzilla:2184487

SSG remediations are now aligned with **configure_openssl_cryptopolicy**

Previously, the SCAP Security Guide (SSG) remediation added the **=** character to the **opensslcnf.config** file. This syntax dit not match the description of the **configure_openssl_cryptopolicy** rule. Consequently, compliance checks might fail after remediations that inserted **.include =** instead of **.include** to **opensslcnf.config**. With this release, the remediation scripts are aligned with the rule description, and SSG remediations that use **configure_openssl_cryptopolicy** no longer fail due to additional **=**.

Bugzilla:2192893

The **postfix_prevent_unrestricted_relay** rule now accepts white spaces around the **=** sign

Previously, the OVAL check of the SCAP rule **xccdf_org.ssgproject.content_rule_postfix_prevent_unrestricted_relay** was too strict and it did not account for **postconf** configuration assignment statements which contained white spaces around the **=** sign. As a consequence, the final report reported this rule as failing even for configurations that technically met the rule's requirements. With this update, the rule was modified so that the check accepts statements with white spaces around the **=** sign. As a result, the final report rule now marks this rule as passing for correct configuration statements.

Bugzilla:2170530

SCAP rules now correctly evaluate whether the **/var/log** and **/var/log/audit** partitions exist

Previously, some SCAP rules relevant to the **/var/log** and **/var/log/audit** partitions were evaluated and remediated even when the appropriate disk partition did not exist. This affected the following rules:

- **mount_option_var_log_audit_nodev**

- **mount_option_var_log_audit_noexec**

- **mount_option_var_log_audit_nosuid**

- **mount_option_var_log_nodev**

- **mount_option_var_log_noexec**

- **mount_option_var_log_nosuid**

As a consequence, these rules were evaluated and incorrectly reported as failing in the final report even when the directories **/var/log** or **/var/log/audit** were not mount points for individual partitions. This update adds an applicability check to determine whether **/var/log** or **/var/log/audit** are mount points for individual partitions. As a consequence, the rules are not evaluated in configurations when the directories are not mount points for individual partitions and the rules are marked as **notapplicable** in the final report.

Bugzilla:2176008

The SCAP rule **accounts_passwords_pam_faillock_interval** now covers new STIG IDs

Previously, the SCAP Security Guide rule **accounts_passwords_pam_faillock_interval** did not cover RHEL-08-020012 and RHEL-08-020013. Consequently, the rule **accounts_passwords_pam_faillock_interval** checked for **faillock** configuration in all of these three files: **/etc/pam.d/password-auth**, **/etc/pam.d/system-auth**, and **/etc/security/faillock.conf**. With this update, the rule now covers STIG IDs RHEL-08-020012 and RHEL-08-020013.

Bugzilla:2209073

Red Hat CVE feeds have been updated

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

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

Bugzilla:2222583

The **wget** utility no longer fails TLS handshake when accessing restricted resources

Previously, when ticket-based session resumption was enabled in TLS, the **wget** utility expected a TLS
session to be resumed even when the server requested the client to re-authenticate to access restricted
resources. This behavior caused **wget** to fail the second TLS handshake. With this update, **wget** properly
initiates a new handshake and the access to restricted resources no longer fails.

Bugzilla:2089817

Settings from **pam_cap** are correctly applied on SELinux-enabled systems

Previously, the SELinux policy did not contain rules for using the **pam_cap** module. As a consequence,
granting login capabilities controlled by **pam_cap** to users in the **/etc/security/capability.conf**
configuration file did not work when the users logged in by using **ssh** or the console. This update adds a
new rule to the policy. As a result, granting capabilities in **/etc/security/capability.conf** now works, and
user capabilities configured with **pam_cap** are taken into account when logging in.

Bugzilla:2172541

The **systemd-fsck-root** service is now correctly labeled on SELinux-enabled systems

Previously, the **/run/fsck** directory was created by the **systemd-fsck-root** service or the **fsck** command
but the SELinux policy did not contain rules for proper labeling of the directory. As a consequence, the
systemd-fsck-root service did not work correctly. With this update, the correct label and file transition
for **/run/fsck** were added to the policy. As a result, the **systemd-fsck-root** service works without
reporting errors.

Bugzilla:2184348[1]

SELinux policy now allows bidirectional communication on D-Bus

Previously, the SELinux policy contained rules to allow only one-way communication between two
domains on the D-Bus message bus system. However, such communication must be allowed in both
directions. This occurred also when the Pacemaker high-availability cluster resource manager executed
the **hostnamectl** or **timedatectl** commands. As a consequence, these commands executed by
Pacemaker timed out without receiving a response on D-Bus because SELinux blocked it. This update to
the SELinux policy allows bidirectional communication on D-Bus. As a result, commands that require
bidirectional communication on D-Bus executed by Pacemaker finish successfully.

Bugzilla:2196524

tangd-keygen now handles non-default umask correctly

Previously, the **tangd-keygen** script did not change file permissions for generated key files.
Consequently, on systems with a default user file-creation mode mask (**umask**) that prevents reading
keys to other users, the **tang-show-keys** command returned the error message **Internal Error 500**
instead of displaying the keys. With this update, **tangd-keygen** sets file permissions for generated key
files, and therefore the script now works correctly on systems with non-default **umask**.

Bugzilla:2188743

Clevis now handles SHA-256 thumbprints

Before this update, the Clevis client did not recognize SHA-256 thumbprints specified through the **thp** configuration option. Consequently, clients did not bind to Tang servers that used SHA-256 thumbprints, and every corresponding **clevis encrypt tang** command reported an error. With this update, Clevis recognizes thumbprints using SHA-256 and handles them correctly. As a result, Clevis clients can bind not only to Tang servers using SHA-1 but also SHA-256 thumbprints.

Bugzilla:2209058

Rsyslog can start even without capabilities

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

Jira:RHELPLAN-160541[1]

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

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

Jira:RHEL-630

fapolicyd service now creates RPM database files with correct ownership

Previously, the **fapolicyd** service created and owned RPM database files in the **/var/lib/rpm/** directory. As a result, other programs were unable to access the files, which resulted in availability control errors. With this update, **fapolicyd** creates the files with correct ownership, and the errors no longer occur.

Jira:RHEL-829

8.3. SOFTWARE MANAGEMENT

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

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

Bugzilla:2122587

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

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

Bugzilla:2170093

YUM now handles proxy=_none_ correctly

You can use the YUM **proxy=_none_** configuration option to prohibit changing proxy settings. Previously, if you set **proxy=_none_** in the main configuration file, YUM detected an error. This update fixes the bug, and YUM now handles **proxy=_none_** correctly.

> **NOTE**
>
> The RHEL 8 YUM **proxy=_none_** configuration is compatible with the YUM configuration in RHEL 7.

Bugzilla:2155713

The needs-restarting plug-in now correctly requires the system restart when a file owned by dbus is updated by zlib

Previously, when you ran the YUM **needs-restarting** plug-in, it did not prompt to restart the system when a file owned by the **dbus** package was updated by the dependent **zlib** package. With this update, the issue has been fixed, and the **needs-restarting** plug-in now displays a message that you must restart **dbus** when **zlib** is updated.

Bugzilla:2092033

8.4. SHELLS AND COMMAND-LINE TOOLS

The which command no longer fails for a long path

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

Bugzilla:2140566

ReaR now supports UEFI Secure Boot with OUTPUT=USB

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

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

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

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

Bugzilla:2233526

ipmievd now recognizes SEL response correctly when a SEL request times out

The **ipmievd** service sends System Event Log (SEL) requests through the /**dev**/**ipmi0** device. Previously, due to a missing ID check of the returned IPMI message, a timed-out request led to incorrect processing of the next request. For example, if the Baseboard Management Controller (BMC) was reset, the SEL request from the **ipmievd** service timed out due to no SEL response. Consequently, **ipmievd** did not work correctly due to a non-corresponding SEL response. As a result, you did not get the correct hardware state, and a large amount of wrong hardware information was output to /**var**/**log**/**messages**. With this fix, **ipmitool** and **ipmievd** now check the ID of the returned IPMI message against the ID of the request and skip non-corresponding SEL requests. **ipmevd** no longer logs incorrect hardware information.

Bugzilla:2224567[1]

8.5. NETWORKING

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

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

Bugzilla:2130727[1]

8.6. BOOT LOADER

grubby now passes arguments to a new kernel correctly

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

Bugzilla:1900829

8.7. FILE SYSTEMS AND STORAGE

multipathd adds the persistent reservation registration key to all paths

Previously, when the **multipathd** daemon started and it recognized a registration key for the persistent reservations on one path of an existing multipath device, not all paths of that device had the registration key. As a consequence, if new paths appeared to a multipath device with persistent reservations while **multipathd** was stopped, persistent reservations were not set up on those. This allowed IO processing on the paths, even if they were supposed to be forbidden by the reservation key.

With this fix, if **multipathd** finds a persistent reservation registration key on any device path, it adds the key to all active paths. As a result, multipath devices now have persistent reservations set up correctly on all the paths, even if path devices first appear while **multipathd** is not running.

Bugzilla:2164871

LUNs are now visible during the OS 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:2213193[1]

8.8. HIGH AVAILABILITY AND CLUSTERS

Pacemaker Designated Controller elections no longer finalized until all pending actions are complete

When a cluster elects a new Designated Controller (DC), all nodes send their current history to the new DC, which saves it to the CIB. As a consequence, if actions were already in progress when a new DC is elected, and the actions finish after the nodes send their current history to the new DC, the actions' results could be lost. With this fix, DC elections are not finalized until all pending actions are complete and no action results are lost.

Bugzilla:2010084

The fence_scsi agent is now able to auto-detect shared lvmlockd devices

Previously, the **fence_scsi** agent did not auto-detect shared **lvmlockd** devices. With this update, **fence_scsi** is able to auto-detect **lvmlockd** devices when the **devices** attribute is not set.

Bugzilla:2187329

Resource stickiness now properly compares against colocation scores

Chained resource colocations are resources colocated with the resource that is colocated with the resource being assigned. Previously, if the original colocation had a finite negative score, and the chained colocation was mandatory, the original resource being assigned could be banned from its node even if resource-stickiness was set to **INFINITY**. With this fix, chained colocations are now taken into account proportionally and stickiness properly compares against colocation scores.

Bugzilla:1632951[1]

The crm_resource command now allows banning or moving a bundle with only a single active replica

Previously, when the **crm_resource** command checked where a bundle with a single replica was active, the command counted both the node where the container was active and the guest node that was created for the container itself. As a result, the **crm_resource** command would not ban or move a bundle with a single active replica. With this fix, the **crm_resource** command now only counts nodes where a bundle's containers are active when determining the number of active replicas.

Bugzilla:1578820

The mysql resource agent now works correctly with promotable clone resources

Previously, the **mysql** resource agent moved cloned resources that were operating in a Master role between nodes, due to promotion scores changing between promoted and non-promoted values. With this fix, a promoted node stays promoted.

Bugzilla:2039692

Unpromoted clone instances no longer restart unnecessarily

Previously, promotable clone instances were assigned in numerical order, with promoted instances first. As a result, if a promoted clone instance needed to start, an unpromoted instance in some cases restarted unexpectedly, because the instance numbers changed. With this fix, roles are considered when assigning instance numbers to nodes and as a result no unnecessary restarts occur.

Bugzilla:1931023

A fence watchdog configured as a second fencing device now fences a node when the first device times out

Previously, when a watchdog fencing device was configured as the second device in a fencing topology, the watchdog timeout would not be considered when calculating the timeout for the fencing operation. As a result, if the first device timed out the fencing operation would time out even though the watchdog would fence the node. With this fix, the watchdog timeout is included in the fencing operation timeout and the fencing operation succeeds if the first device times out.

Bugzilla:2168633

Location constraints with rules no longer displayed when listing is grouped by nodes

Location constraints with rules cannot have a node assigned. Previously, when you grouped the listing by nodes, location constraints with rules were displayed under an empty node. With this fix, the location constraints with rules are no longer displayed and a warning is given indicating that constraints with rules are not displayed.

Bugzilla:2166294

pcs command to update multipath SCSI devices now works correctly

Due to changes in the Pacemaker CIB file, the **pcs stonith update-scsi-devices** command stopped working as designed, causing an unwanted restart of some cluster resources. With this fix, this command works correctly and updates SCSI devices without requiring a restart of other cluster resources running on the same node.

Bugzilla:2179010

Memory footprint of `pcsd-ruby` daemon now reduced when `pscd` Web UI is open

Previously, when the **pcsd** Web UI was open, memory usage of the **pcsd-ruby** daemon increased steadily over the course of several hours. With this fix, the web server that runs in the **pcsd-ruby** daemon now periodically performs a graceful restart. This frees the allocated memory and reduces the memory footprint.

Bugzilla:2189958[1]

The `azure-events-az` resource agent no longer produces an error with Pacemaker 2.1 and later

The **azure-events-az** resource agent executes the **crm_simulate -Ls** command and parses the output. With Pacemaker 2.1 and later, the output of the **crm_simulate** command no longer contains the text **Transition Summary:**, which resulted in an error. With this fix, the agent no longer yields an error when this text is missing.

Bugzilla:2181019

8.9. COMPILERS AND DEVELOPMENT TOOLS

systemtap scripts using guru mode now compile more quickly

The **systemtap** guru mode liveness analysis uses the **dyninst** library to parse binaries. Newer kernels enable mitigation code with **CONFIG_RETPOLINE=y**, replacing traditional RET instructions, with jumps to a thunk. As a consequence, binary analysis took a much longer time due to the liveness analysis needing to examine all additional edges of the control flow graph introduced by the jumps to the thunk.

With this update, **systemtap** disables liveness analysis when the kernel code is using thunks and, as a result, **systemtap** scripts using guru mode compile more quickly.

Bugzilla:2126805

eu-addr2line -C now correctly recognizes other arguments

Previously, when you used the **-C** argument in **eu-addr2line** command from **elfutils**, the following single character argument disappeared. Consequently, the **eu-addr2line -Ci** command behaved the same way as **eu-addr2line -C** while **eu-addr2line -iC** worked as expected. This bug has been fixed, and **eu-addr2line -Ci** now recognizes both arguments.

Bugzilla:2236183

eu-addr2line -i now correctly handles code compiled with GCC link-time optimization

Previously, the **dwarf_getscopes** function from the **libdw** library included in **elfutils** was unable to find an abstract origin definition of a function that was compiled with GCC link-time optimization. Consequently, when you used the **-i** argument in the **eu-addr2line** command, **eu-addr2line** was unable to show inline functions for code compiled with **gcc -flto**. With this update, the **libdw dwarf_getscopes** function looks in the correct compile unit for the inlined scope, and **eu-addr2line -i** works as expected.

Bugzilla:2162495

8.10. IDENTITY MANAGEMENT

SSSD now uses sAMAccountName when evaluating GPO-based access control

Previously, if **ldap_user_name** was set to a value other than **sAMAccountName** on an AD client, GPO-based access control failed. With this update, SSSD now always uses **sAMAccountName** when evaluating GPO-based access control. Even if **ldap_user_name** is set to a value different from **sAMAccountName** on an AD client, GPO-based access control now works correctly.

Jira:SSSD-6107

SSSD now handles duplicate attributes in the user_attributes option when retrieving users

Previously, if **sssd.conf** contained duplicate attributes in the **user_attributes** option, SSSD did not handle these duplicates correctly. As a consequence, users with those attributes could not be retrieved. With this update, SSSD now handles duplicates correctly. As a result, users with duplicate attributes can now be retrieved.

Jira:SSSD-6177

Changing a security parameter now works correctly

Previously, when you changed a security parameter by using the **dsconf** *instance_name* **security set** command, the operation failed with the error:

> Name 'log' is not defined

With this update, the security parameter change works as expected.

Directory Server now calculates the dtablesize based on the maximum number of opened descriptors

Previously, an administrator could set the connection table size manually by using the **nsslapd-conntablesize** configuration parameter. Consequently, when the connection table size was set too low, it affected the number of connections the server was able to support. With this update, Directory Server now calculates the size of the connection table dynamically effectively resolving the issue with too small connection table size. In addition, you no longer need to manually change the connection table size.

The dsctl healthcheck command now uses the password storage scheme PBKDF2-SHA512 by default

Previously, the **dsctl healthcheck** command used **SSHA512** password storage scheme by default. Consequently, the command reported a warning because it did not detect the new password storage scheme **PBKDF2-SHA512**. With this update, the **dsctl healthcheck** command now uses **PBKDF2-SHA512** password storage scheme by default and no warnings occur.

Paged searches from a regular user now do not impact performance

Previously, when Directory Server was under the search load, paged searches from a regular user could impact the server performance because a lock conflicted with the thread that polls for network events. In addition, if a network issue occurred while sending the page search, the whole server was unresponsive until the **nsslapd-iotimeout** parameter expired. With this update, the lock was split into several parts to avoid the contention with the network events. As a result, no performance impact during paged searches from a regular user.

You can now enable and disable ciphers in Directory Server as expected

Previously, when you tried to enable or disable specific ciphers in addition to default ciphers by using the web console, the server enabled or disabled only the specific ciphers and logged an error similar to the following:

> Security Initialization - SSL alert: Failed to set SSL cipher preference information: invalid ciphers <default,+cipher_name>: format is +cipher1,-cipher2... (Netscape Portable Runtime error 0 - no error)

Currently, the network security services (NSS) do not support handling default ciphers and specific ciphers at the same time. As a result, Directory Server can enable or disable either specific ciphers or default ciphers. With this update, when you set the default ciphers, the web console now prompts that **Allow Specific Ciphers** and **Deny Specific Ciphers** fields will be cleared.

Deleting the IdM admin user is now no longer permitted

Previously, nothing prevented you from deleting the Identity Management (IdM) **admin** user if you were a member of the **admins** group. The absence of the **admin** user causes the trust between IdM and Active Directory (AD) to stop functioning correctly. With this update, you can no longer delete the **admin** user. As a result, the IdM–AD trust works correctly.

Bugzilla:1821181

IdM clients correctly retrieve information for trusted AD users when their names contain mixed case characters

Previously, if you attempted a user lookup or authentication of a user, and that trusted Active Directory (AD) user contained mixed case characters in their names and they were configured with overrides in IdM, an error was returned preventing users from accessing IdM resources.

With the release of RHBA-2023:4525, a case-sensitive comparison is replaced with a case-insensitive comparison that ignores the case of a character. As a result, IdM clients can now lookup users of an AD trusted domain, even if their usernames contain mixed case characters and they are configured with overrides in IdM.

Jira:SSSD-6096

8.11. GRAPHICS INFRASTRUCTURES

The installer no longer freezes on servers with ASPEED 2600

Previously, the graphical RHEL 8.8 installer became unresponsive with a black screen when you started the installer on a server with the ASPEED 2600 On System Management Chipset. Consequently, you could not install RHEL 8.8 on the server.

With this release, the problem has been fixed. As a result, the installation now proceeds as expected with ASPEED 2600.

Bugzilla:2189645[1]

8.12. THE WEB CONSOLE

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

In RHEL 8.8, due to a bug in the code for determining whether or not the user was adding a Tang key to the root file system, the binding process in the web console crashed when there was no file system on the LUKS container at all. Because the web console displayed the error message **TypeError: Qe(…) is undefined** after you had clicked the **Trust key** button in the **Verify key** dialog, you had to perform all the required steps in the command-line interface in the described scenario.

With this update, the web console correctly handles additions of Tang keys to root file systems. As a result, the web console finishes all binding steps required for the automated unlocking of LUKS-encrypted volumes using Network-Bound Disk Encryption (NBDE) in various scenarios.

Bugzilla:2212350

VNC console now works at most resolutions

Previously, when using the Virtual Network Computing (VNC) console under certain display resolutions, a mouse offset problem was present or only a part of the interface was visible. Consequently, using the VNC console was not possible.

With this update, the problem has been fixed and the VNC console works correctly at most resolutions, with the exception of ultra high resolutions, such as 3840x2160.

Note that a small offset between the recorded and displayed positions of the cursor might still be present. However, this does not significantly impact the usability of the VNC console.

Bugzilla:2030836

8.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The storage role can now resize the mounted file systems without unmounting

Previously, the **storage** role was unable to resize mounted devices, even if the file system supported online resizing. As a consequence, the **storage** role unmounted all file systems prior to resizing, which failed for file systems that were in use, for example, while resizing the / directory of the running system.

With this update, the **storage** role now supports resizing mounted file systems that support online resizing such as XFS and Ext4. As a result, the mounted file systems can now be resized without unmounting them.

Bugzilla:2168738

The certificate RHEL System Role now checks for the certificate key size when determining whether to perform a new certificate request

Previously, the **certificate** RHEL System Role did not check the key size of a certificate when evaluating whether to request a new certificate. As a consequence, the role sometimes did not issue new certificate requests in cases where it should. With this update, **certificate** now checks the **key_size** parameter to determine if a new certificate request should be performed.

Bugzilla:2186057

Insights tags created by using the rhc role are now applied correctly

Previously, when you created Insights tags by using the **rhc** role, tags were not stored in the correct file. Consequently, tags were not sent to Insights and as a result they were not applied to the systems in the Insights inventory.

With this fix, tags are stored correctly and applied to the systems present in the Insights inventory.

Bugzilla:2209441

The firewall RHEL System Role on RHEL 7 no longer attempts to install non-existent Python packages

Previously, when the **firewall** role on RHEL 7 was called from another role, and that role was using **python3**, the **firewall** role attempted to install the **python3-firewall** library for that version of Python. However, that library is not available in RHEL 7. Consequently, the **python3-firewall** library was not found, and you received the following error message:

> No package matching 'python3-firewall' found available, installed or updated

With this update, the **firewall** role does not attempt to install the **python-firewall** or **python3-firewall** library. As a result, the **firewall** role does not fail on RHEL 7 when **python3** is installed on the managed node.

Bugzilla:2216521

Failure to remove data from member disks before creation no longer persists

Previously, when creating RAID volumes, the system did not effectively eliminate existing data from member disks before forming the RAID volume. With this update, RAID volumes remove any per-existing data from member disks as needed.

Bugzilla:2224094

The podman_registries_conf variable now configures unqualified-search-registries field correctly

Previously, after configuring the **podman_registries_conf** variable, the **podman** RHEL System Role failed. Consequently, **unqualified-search-registries = ["registry.access.redhat.com"]** setting was not generated in the **/etc/containers/registries.conf.d/50-systemroles.conf** file. With this update, this problem has been fixed.

Bugzilla:2226077

raid_chunk_size parameter no longer returns an error message

Previously, **raid_chunk_size** attribute was not allowed for RAID pools and volumes. With this update, you can now configure the **raid_chunk_size** attribute for RAID pools and volumes without encountering any restrictions.

Bugzilla:2193057

Running the firewall RHEL System Role in check mode with non-existent services no longer fails

Previously, running the **firewall** role in check mode with non-existent services would fail. This fix implements better compliance with Ansible best practices for check mode. As a result, non-existent services being enabled or disabled no longer fails the role in check mode. Instead, a warning prompts you to confirm that the service is defined in a previous playbook.

Bugzilla:2222433

The kdump role adds authorized_keys idempotently

Previously, the task to add **authorized_key** added an extra newline character every time. Consequently the role was not acting idempotent. With this fix, adding a new **authorized_key** works correctly and adds only a single key value idempotently.

Bugzilla:2232391

The kdump system role does not fail if authorized_keys are missing

Previously, the **kdump** system role failed to add **SSH** authorized keys if the user defined in the **kdump_ssh_user** variable did not have access to the **.ssh** directory in the **home** directory or an empty **.ssh/authorized_keys** file. With this fix, the **kdump** system role now correctly adds authorized keys to the **SSH** configuration. As a result, the key based authentication works reliably in the described scenario.

Bugzilla:2232392

The firewall RHEL System Role correctly reports changes when using previous: replaced in check mode

Previously, the **firewall** role was not checking whether any files would be changed when using the **previous: replaced** parameter in check mode. As a consequence, the role gave an error about undefined variables. This fix adds new check variables to the check mode to assess whether any files would be

changed by the **previous: replaced** parameter. The check for the **firewalld.conf** file assesses the **rpm** database to determine whether the file has been changed from the version shipped in the package. As a result, the **firewall** role now correctly reports changes when using the **previous: replaced** parameter.

Jira:RHEL-899[1]

Enabling **kdump** for system role requires using the **failure_action** configuration parameter on RHEL 9 and later versions

Previously, using the **default** option during **kdump** configuration was not successful and printed the following warning in logs:

> kdump: warning: option 'default' was renamed 'failure_action' and will be removed in the future. please update /etc/kdump.conf to use option 'failure_action' instead.

Consequently, the role did not enable **kdump** successfully if **default** option was used. This update fixes the problem and you can configure kernel dump parameters on multiple systems by using the **failure_action** parameter. As a result, enabling **kdump** works successfully in the described scenario.

Jira:RHEL-907[1]

The **firewall** RHEL System Role correctly reports changes when assigning zones to Network Manager interfaces

Previously, the Network Manager interface assignment reported changes when no changes were present. With this fix, the **try_set_zone_of_interface** module in the file **library/firewall_lib.py** returns a second value, which denotes whether the interface's zone was changed. As a result, the module now correctly reports changes when assigning zones to interfaces handled by Network Manager.

Jira:RHEL-918[1]

The **kdump** role successfully updates .ssh/authorized_keys for **kdump_ssh_server** authentication

Previously, the **.ssh** directory was not accessible by the **kdump** role to securely authenticate users to log into **kdump_ssh_server**. As a consequence, the **kdump** role did not update the **.ssh/authorized_keys** file and the SSH mechanism to verify the **kdump_ssh_server** failed. This update fixes the problem. As a result the **kdump_ssh_user** authentication on **kdump_ssh_server** works reliably.

Jira:RHEL-1398[1]

The **previous: replaced** parameter of the **firewall** System Role now overrides the previous configuration without deleting it

Previously, if you added the **previous: replaced** parameter to the variable list, the **firewall** System Role removed all existing user-defined settings and reset **firewalld** to the default settings. This fix uses the fallback configuration in **firewalld**, which was introduced in the EL7 release, to retain the previous configuration. As a result, when you use the **previous: replaced** parameter in the variable list, the **firewall.conf** configuration file is not deleted on reset, but the file and comments in the file are retained.

Jira:RHEL-1496[1]

The **kdump** role adds multiple keys to authorized_keys idempotently

Previously, adding multiple SSH keys to the **authorized_keys** file at the same time replaced the key

value of one host by another. This update fixes the problem by using the **lineinfile** module to manage the **authorized_keys** file. **lineinfile** iterates the tasks in sequence, checking for an existing key and writing the new key in one atomic operation on a single host at one time. As a result, adding SSH keys on multiple hosts works correctly, and does not replace the key value from another host.

Note: Use the **serial: 1** play serial keyword at play level to control the number of hosts executing at one time.

Jira:RHEL-1500[1]

8.14. VIRTUALIZATION

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

CHAPTER 9. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 8.9.

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

9.1. 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

9.2. NETWORKING

AF_XDP available as a Technology Preview

Address Family eXpress Data Path (**AF_XDP**) socket is designed for high-performance packet processing. It accompanies **XDP** and grants efficient redirection of programmatically selected packets to user space applications for further processing.

Bugzilla:1633143[1]

XDP features that are available as Technology Preview

Red Hat provides the usage of the following eXpress Data Path (XDP) features as unsupported Technology Preview:

- Loading XDP programs on architectures other than AMD and Intel 64-bit. Note that the **libxdp** library is not available for architectures other than AMD and Intel 64-bit.

- The XDP hardware offloading.

Bugzilla:1889737

Multi-protocol Label Switching for TC available as a Technology Preview

The Multi-protocol Label Switching (MPLS) is an in-kernel data-forwarding mechanism to route traffic flow across enterprise networks. In an MPLS network, the router that receives packets decides the further route of the packets based on the labels attached to the packet. With the usage of labels, the MPLS network has the ability to handle packets with particular characteristics. For example, you can add **tc filters** for managing packets received from specific ports or carrying specific types of traffic, in a consistent way.

After packets enter the enterprise network, MPLS routers perform multiple operations on the packets, such as **push** to add a label, **swap** to update a label, and **pop** to remove a label. MPLS allows defining actions locally based on one or multiple labels in RHEL. You can configure routers and set traffic control (**tc**) filters to take appropriate actions on the packets based on the MPLS label stack entry (**lse**) elements, such as **label**, **traffic class**, **bottom of stack**, and **time to live**.

For example, the following command adds a filter to the *enp0s1* network interface to match incoming packets having the first label *12323* and the second label *45832*. On matching packets, the following actions are taken:

- the first MPLS TTL is decremented (packet is dropped if TTL reaches 0)

- the first MPLS label is changed to *549386*

- the resulting packet is transmitted over *enp0s2*, with destination MAC address *00:00:5E:00:53:01* and source MAC address *00:00:5E:00:53:02*

```
# tc filter add dev enp0s1 ingress protocol mpls_uc flower mpls lse depth 1 label 12323 lse
depth 2 label 45832 \
action mpls dec_ttl pipe \
action mpls modify label 549386 pipe \
action pedit ex munge eth dst set 00:00:5E:00:53:01 pipe \
action pedit ex munge eth src set 00:00:5E:00:53:02 pipe \
action mirred egress redirect dev enp0s2
```

Bugzilla:1814836[1], Bugzilla:1856415

act_mpls module available as a Technology Preview

The **act_mpls** module is now available in the **kernel-modules-extra** rpm as a Technology Preview. The module allows the application of Multiprotocol Label Switching (MPLS) actions with Traffic Control (TC) filters, for example, push and pop MPLS label stack entries with TC filters. The module also allows the Label, Traffic Class, Bottom of Stack, and Time to Live fields to be set independently.

Bugzilla:1839311[1]

The **systemd-resolved** service is now 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, even if the **systemd** package provides **systemd-resolved**, this service is an unsupported Technology Preview.

Bugzilla:1906489

9.3. KERNEL

Soft-RoCE available as a Technology Preview

Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) is a network protocol that implements RDMA over Ethernet. Soft-RoCE is the software implementation of RoCE which maintains two protocol versions, RoCE v1 and RoCE v2. The Soft-RoCE driver, **rdma_rxe**, is available as an unsupported Technology Preview in RHEL 8.

Bugzilla:1605216[1]

eBPF available as a Technology Preview

Extended Berkeley Packet Filter (eBPF) is an in-kernel virtual machine that allows code execution in the kernel space, in the restricted sandbox environment with access to a limited set of functions.

The virtual machine includes a new system call **bpf()**, which enables creating various types of maps, and also allows to load programs in a special assembly-like code. The code is then loaded to the kernel and translated to the native machine code with just-in-time compilation. Note that the **bpf()** syscall can be successfully used only by a user with the **CAP_SYS_ADMIN** capability, such as the root user. See the **bpf(2)** manual page for more information.

The loaded programs can be attached onto a variety of points (sockets, tracepoints, packet reception) to receive and process data.

There are numerous components shipped by Red Hat that utilize the **eBPF** virtual machine. Each component is in a different development phase. All components are available as a Technology Preview, unless a specific component is indicated as supported.

The following notable **eBPF** components are currently available as a Technology Preview:

- **AF_XDP**, a socket for connecting the **eXpress Data Path (XDP)** path to user space for applications that prioritize packet processing performance.

Bugzilla:1559616[1]

The **kexec** fast reboot feature is available as a Technology Preview

The **kexec** fast reboot feature continues to be available as a Technology Preview. The **kexec** fast reboot significantly speeds the boot process as you can boot directly into the second kernel without passing through the Basic Input/Output System (BIOS) or firmware first. To use this feature:

1. Load the **kexec** kernel manually.

2. Reboot for changes to take effect.

Note that the **kexec** fast reboot capability is available with a limited scope of support on RHEL 9 and later releases.

Bugzilla:1769727

The Intel data streaming accelerator driver for kernel is available as a Technology Preview

The Intel data streaming accelerator driver (IDXD) for the kernel is currently available as a Technology Preview. It is an Intel CPU integrated accelerator and includes a shared work queue with process address space ID (pasid) submission and shared virtual memory (SVM).

Bugzilla:1837187[1]

The **accel-config** package available as a Technology Preview

The **accel-config** package is now available on Intel **EM64T** and **AMD64** architectures as a Technology Preview. This package helps in controlling and configuring data-streaming accelerator (DSA) sub-system in the Linux Kernel. Also, it configures devices through **sysfs** (pseudo-filesystem), saves and loads the configuration in the **json** format.

Bugzilla:1843266[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]

9.4. FILE SYSTEMS AND STORAGE

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

In Red Hat Enterprise Linux 8, the file system DAX is available as a Technology Preview. DAX provides a 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 file system that provides the capability of DAX must be created on the NVDIMM(s). Also, the file system must be mounted with the **dax** mount option. Then, a **mmap** of a file on the dax-mounted file system results in a direct mapping of storage into the application's address space.

Bugzilla:1627455[1]

OverlayFS

OverlayFS is a type of union file system. It enables you to overlay one file system on top of another. Changes are recorded in the upper file system, while the lower file system remains unmodified. This allows multiple users to share a file-system image, such as a container or a DVD-ROM, where the base image is on read-only media.

OverlayFS remains a Technology Preview under most circumstances. As such, the kernel logs warnings when this technology is activated.

Full support is available for OverlayFS when used with supported container engines (**podman**, **cri-o**, or **buildah**) under the following restrictions:

- OverlayFS is supported for use only as a container engine graph driver or other specialized use cases, such as squashed **kdump** initramfs. Its use is supported primarily for container COW content, not for persistent storage. You must place any persistent storage on non-OverlayFS volumes. You can use only the default container engine configuration: one level of overlay, one lowerdir, and both lower and upper levels are on the same file system.

- Only XFS is currently supported for use as a lower layer file system.

Additionally, the following rules and limitations apply to using OverlayFS:

- The OverlayFS kernel ABI and user-space behavior are not considered stable, and might change in future updates.

- OverlayFS provides a restricted set of the POSIX standards. Test your application thoroughly before deploying it with OverlayFS. The following cases are not POSIX-compliant:

  - Lower files opened with **O_RDONLY** do not receive **st_atime** updates when the files are read.

- Lower files opened with **O_RDONLY**, then mapped with **MAP_SHARED** are inconsistent with subsequent modification.

- Fully compliant **st_ino** or **d_ino** values are not enabled by default on RHEL 8, but you can enable full POSIX compliance for them with a module option or mount option.
  To get consistent inode numbering, use the **xino=on** mount option.

  You can also use the **redirect_dir=on** and **index=on** options to improve POSIX compliance. These two options make the format of the upper layer incompatible with an overlay without these options. That is, you might get unexpected results or errors if you create an overlay with **redirect_dir=on** or **index=on**, unmount the overlay, then mount the overlay without these options.

- To determine whether an existing XFS file system is eligible for use as an overlay, use the following command and see if the **ftype=1** option is enabled:

  ```
  # xfs_info /mount-point | grep ftype
  ```

- SELinux security labels are enabled by default in all supported container engines with OverlayFS.

- Several known issues are associated with OverlayFS in this release. For details, see *Non-standard behavior* in the Linux kernel documentation.

For more information about OverlayFS, see the Linux kernel documentation.

Bugzilla:1690207[1]

Stratis is now available as a Technology Preview

Stratis is a new local storage manager, which provides managed file systems on top of pools of storage with additional features. It is provided as a Technology Preview.

With Stratis, you can perform the following storage tasks:

- Manage snapshots and thin provisioning

- Automatically grow file system sizes as needed

- Maintain file systems

To administer Stratis storage, use the **stratis** utility, which communicates with the **stratisd** background service. For more information, see the Setting up Stratis file systems documentation.

RHEL 8.5 updated Stratis to version 2.4.2. For more information, see the Stratis 2.4.2 Release Notes.

Jira:RHELPLAN-1212[1]

NVMe/TCP host is available as a Technology Preview

Accessing and sharing Nonvolatile Memory Express (NVMe) storage over TCP/IP networks (NVMe/TCP) and its corresponding **nvme_tcp.ko** kernel module has been added as a Technology Preview. The use of NVMe/TCP as a host is manageable with tools provided by the **nvme-cli** package. The NVMe/TCP host Technology Preview is included only for testing purposes and is not currently planned for full support.

Bugzilla:1696451[1]

Setting up a Samba server on an IdM domain member is provided as a Technology Preview

With this update, you can now set up a Samba server on an Identity Management (IdM) domain member. The new **ipa-client-samba** utility provided by the same-named package adds a Samba-specific Kerberos service principal to IdM and prepares the IdM client. For example, the utility creates the **/etc/samba/smb.conf** with the ID mapping configuration for the **sss** ID mapping back end. As a result, administrators can now set up Samba on an IdM domain member.

Due to IdM Trust Controllers not supporting the Global Catalog Service, AD-enrolled Windows hosts cannot find IdM users and groups in Windows. Additionally, IdM Trust Controllers do not support resolving IdM groups using the Distributed Computing Environment / Remote Procedure Calls (DCE/RPC) protocols. As a consequence, AD users can only access the Samba shares and printers from IdM clients.

For details, see Setting up Samba on an IdM domain member .

Jira:RHELPLAN-13195[1]

9.5. HIGH AVAILABILITY AND CLUSTERS

Pacemaker **podman** bundles available as a Technology Preview

Pacemaker container bundles now run on Podman, with the container bundle feature being available as a Technology Preview. There is one exception to this feature being Technology Preview: Red Hat fully supports the use of Pacemaker bundles for Red Hat OpenStack.

Bugzilla:1619620[1]

Heuristics in **corosync-qdevice** available as a Technology Preview

Heuristics are a set of commands executed locally on startup, cluster membership change, successful connect to **corosync-qnetd**, and, optionally, on a periodic basis. When all commands finish successfully on time (their return error code is zero), heuristics have passed; otherwise, they have failed. The heuristics result is sent to **corosync-qnetd** where it is used in calculations to determine which partition should be quorate.

Bugzilla:1784200

New **fence-agents-heuristics-ping** fence agent

As a Technology Preview, Pacemaker now provides the **fence_heuristics_ping** agent. This agent aims to open a class of experimental fence agents that do no actual fencing by themselves but instead exploit the behavior of fencing levels in a new way.

If the heuristics agent is configured on the same fencing level as the fence agent that does the actual fencing but is configured before that agent in sequence, fencing issues an **off** action on the heuristics agent before it attempts to do so on the agent that does the fencing. If the heuristics agent gives a negative result for the **off** action it is already clear that the fencing level is not going to succeed, causing Pacemaker fencing to skip the step of issuing the **off** action on the agent that does the fencing. A heuristics agent can exploit this behavior to prevent the agent that does the actual fencing from fencing a node under certain conditions.

A user might want to use this agent, especially in a two-node cluster, when it would not make sense for a node to fence the peer if it can know beforehand that it would not be able to take over the services properly. For example, it might not make sense for a node to take over services if it has problems reaching the networking uplink, making the services unreachable to clients, a situation which a ping to a router might detect in that case.

Bugzilla:1775847[1]

9.6. IDENTITY MANAGEMENT

Identity Management JSON-RPC API available as Technology Preview

An API is available for Identity Management (IdM). To view the API, IdM also provides an API browser as a Technology Preview.

Previously, the IdM API was enhanced to enable multiple versions of API commands. These enhancements could change the behavior of a command in an incompatible way. Users are now able to continue using existing tools and scripts even if the IdM API changes. This enables:

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

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

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

For details on using the API, see Using the Identity Management API to Communicate with the IdM Server (TECHNOLOGY PREVIEW).

Bugzilla:1664719

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:1664718

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.
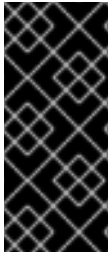
> **IMPORTANT**
>
> It is recommended to enable ACME only in an IdM deployment where all servers are running RHEL 8.4 or later. Earlier RHEL versions do not include the ACME service, which can cause problems in mixed-version deployments. For example, a CA server without ACME can cause client connections 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:1628987[1]

sssd-idp sub-package available as a Technology Preview

The **sssd-idp** sub-package for SSSD contains the **oidc_child** and krb5 **idp** plugins, which are client-side components that perform OAuth2 authentication against Identity Management (IdM) servers. This feature is available only with IdM servers on RHEL 8.7 and later.

Bugzilla:2065692

SSSD internal krb5 idp plugin available as a Technology Preview

The SSSD krb5 **idp** plugin allows you to authenticate against an external identity provider (IdP) using the OAuth2 protocol. This feature is available only with IdM servers on RHEL 8.7 and later.

Bugzilla:2056483

RHEL IdM allows delegating user authentication to external identity providers as a Technology Preview

As a Technology Preview in RHEL IdM, you can now associate users with external identity providers (IdP) that support the OAuth 2 device authorization flow. When these users authenticate with the SSSD version available in RHEL 8.7 or later, they receive RHEL IdM single sign-on capabilities with Kerberos tickets after performing authentication and authorization at the external IdP.

Notable features include:

- Adding, modifying, and deleting references to external IdPs with **ipa idp-*** commands

- Enabling IdP authentication for users with the **ipa user-mod --user-auth-type=idp** command

For additional information, see Using external identity providers to authenticate to IdM .

Bugzilla:2101770

9.7. 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], Bugzilla:1667225, Bugzilla:1724302, Bugzilla:1667516

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]

9.8. GRAPHICS INFRASTRUCTURES

VNC remote console available as a Technology Preview for the 64-bit ARM architecture

On the 64-bit ARM architecture, the Virtual Network Computing (VNC) remote console is available as a Technology Preview. Note that the rest of the graphics stack is currently unverified for the 64-bit ARM architecture.

Bugzilla:1698565[1]

9.9. VIRTUALIZATION

KVM virtualization is usable in RHEL 8 Hyper-V virtual machines

As a Technology Preview, nested KVM virtualization can now be used on the Microsoft Hyper-V hypervisor. As a result, you can create virtual machines on a RHEL 8 guest system running on a Hyper-V host.

Note that currently, this feature only works on Intel and AMD systems. In addition, nested virtualization is in some cases not enabled by default on Hyper-V. To enable it, see the following Microsoft documentation:

https://docs.microsoft.com/en-us/virtualization/hyper-v-on-windows/user-guide/nested-virtualization

Bugzilla:1519039[1]

AMD SEV and SEV-ES for KVM virtual machines

As a Technology Preview, RHEL 8 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 8 includes SEV and SEV-ES encryption, but not the SEV and SEV-ES security attestation.

Bugzilla:1501618[1], Jira:RHELPLAN-7677, Bugzilla:1501607

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 only selected Intel GPUs are compatible with the vGPU feature.

In addition, it is possible to enable a VNC console operated by Intel vGPU. By enabling it, users can connect to a VNC console of the VM and see the VM's desktop hosted by Intel vGPU. However, this currently only works for RHEL guest operating systems.

Note that this feature is deprecated and will be removed entirely in a future RHEL major release.

Bugzilla:1528684[1]

Creating nested virtual machines

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

Jira:RHELPLAN-14047[1], Jira:RHELPLAN-24437

Technology Preview: Select Intel network adapters now provide SR-IOV in RHEL guests on Hyper-V

As a Technology Preview, Red Hat Enterprise Linux guest operating systems running on a Hyper-V hypervisor can now use the single-root I/O virtualization (SR-IOV) feature for Intel network adapters that are supported by the **ixgbevf** and **iavf** drivers. This feature is enabled when the following conditions are met:

- SR-IOV support is enabled for the network interface controller (NIC)

- SR-IOV support is enabled for the virtual NIC

- SR-IOV support is enabled for the virtual switch

- The virtual function (VF) from the NIC is attached to the virtual machine

The feature is currently provided with Microsoft Windows Server 2016 and later.

Bugzilla:1348508[1]

Intel TDX in RHEL guests

As a Technology Preview, the Intel Trust Domain Extension (TDX) feature can now be used in RHEL 8.8 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:1836977[1]

Sharing files between hosts and VMs using virtiofs

As a Technology Preview, RHEL 8 now provides the virtio file system (**virtiofs**). Using **virtiofs**, you can efficiently share files between your host system and its virtual machines (VM).

Bugzilla:1741615[1]

9.10. RHEL IN CLOUD ENVIRONMENTS

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

With the updated RHEL kernel, you can now create and run confidential virtual machines (VMs) on Microsoft Azure as a Technology Preview. However, it is not yet possible to encrypt RHEL confidential VM images during boot on Azure.

Jira:RHELPLAN-122316[1]

9.11. CONTAINERS

SQLite database backend for Podman is available as a Technology Preview

Beginning with Podman v4.6, the SQLite database backend for Podman is available as a Technology Preview. To set the database backend to SQLite, add the **database_backend = "sqlite"** option in the **/etc/containers/containers.conf** configuration file. Run the **podman system reset** command to reset storage back to the initial state before you switch to the SQLite database backend. Note that you have to recreate all containers and pods. The SQLite database guarantees good stability and consistency. Other databases in the containers stack will be moved to SQLite as well. The BoltDB remains the default database backend.

Jira:RHELPLAN-154428[1]

The **podman-machine** command is unsupported

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

Jira:RHELDOCS-16861[1]

CHAPTER 10. DEPRECATED FUNCTIONALITY

This part provides an overview of functionality that has been *deprecated* in Red Hat Enterprise Linux 8.

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 8. 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 7 but has been *removed* in RHEL 8, see Considerations in adopting RHEL 8 .

10.1. INSTALLER AND IMAGE CREATION

Several Kickstart commands and options have been deprecated

Using the following commands and options in RHEL 8 Kickstart files will print a warning in the logs:

- **auth** or **authconfig**

- **device**

- **deviceprobe**

- **dmraid**

- **install**

- **lilo**

- **lilocheck**

- **mouse**

- **multipath**

- **bootloader --upgrade**

- **ignoredisk --interactive**

- **partition --active**

- **reboot --kexec**

Where only specific options are listed, the base command and its other options are still available and not deprecated.

For more details and related changes in Kickstart, see the Kickstart changes section of the *Considerations in adopting RHEL 8* document.

Bugzilla:1642765[1]

The --interactive option of the ignoredisk Kickstart command has been deprecated

Using the **--interactive option** in future releases of Red Hat Enterprise Linux will result in a fatal installation error. It is recommended that you modify your Kickstart file to remove the option.

Bugzilla:1637872[1]

The Kickstart autostep command has been deprecated

The **autostep** command has been deprecated. The related section about this command has been removed from the RHEL 8 documentation .

Bugzilla:1904251[1]

10.2. SECURITY

NSS SEED ciphers are deprecated

The Mozilla Network Security Services (**NSS**) library will not support TLS cipher suites that use a SEED cipher in a future release. To ensure smooth transition of deployments that rely on SEED ciphers when NSS removes support, Red Hat recommends enabling support for other cipher suites.

Note that SEED ciphers are already disabled by default in RHEL.

Bugzilla:1817533

TLS 1.0 and TLS 1.1 are deprecated

The TLS 1.0 and TLS 1.1 protocols are disabled in the **DEFAULT** system-wide cryptographic policy level. If your scenario, for example, a video conferencing application in the Firefox web browser, requires using the deprecated protocols, switch the system-wide cryptographic policy to the **LEGACY** level:

```
# update-crypto-policies --set LEGACY
```

For more information, see the Strong crypto defaults in RHEL 8 and deprecation of weak crypto algorithms Knowledgebase article on the Red Hat Customer Portal and the **update-crypto-policies(8)** man page.

Bugzilla:1660839

DSA is deprecated in RHEL 8

The Digital Signature Algorithm (DSA) is considered deprecated in Red Hat Enterprise Linux 8. Authentication mechanisms that depend on DSA keys do not work in the default configuration. Note that **OpenSSH** clients do not accept DSA host keys even in the **LEGACY** system-wide cryptographic policy level.

Bugzilla:1646541[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:2054741

SSL2 Client Hello has been deprecated in NSS

The Transport Layer Security (**TLS**) protocol version 1.2 and earlier allow to start a negotiation with a **Client Hello** message formatted in a way that is backward compatible with the Secure Sockets Layer (**SSL**) protocol version 2. Support for this feature in the Network Security Services (**NSS**) library has been deprecated and it is disabled by default.

Applications that require support for this feature need to use the new **SSL_ENABLE_V2_COMPATIBLE_HELLO** API to enable it. Support for this feature may be removed completely in future releases of Red Hat Enterprise Linux 8.

Bugzilla:1645153[1]

NTLM and Krb4 are deprecated in Cyrus SASL

The NTLM and Kerberos 4 authentication protocols have been deprecated and might be removed in a future major version of RHEL. These protocols are no longer considered secure and have already been removed from upstream implementations.

Jira:RHELDOCS-17380[1]

Runtime disabling SELinux using /etc/selinux/config is now deprecated

Runtime disabling SELinux using the **SELINUX=disabled** option in the **/etc/selinux/config** file has been deprecated. In RHEL 9, when you disable SELinux only through **/etc/selinux/config**, the system starts with SELinux enabled but with no policy loaded.

If your scenario really requires to completely disable SELinux, Red Hat recommends disabling SELinux by adding the **selinux=0** parameter to the kernel command line as described in the Changing SELinux modes at boot time section of the Using SELinux title.

Bugzilla:1932222

The ipa SELinux module removed from selinux-policy

The **ipa** SELinux module has been removed from the **selinux-policy** package because it is no longer maintained. The functionality is now included in the **ipa-selinux** subpackage.

If your scenario requires the use of types or interfaces from the **ipa** module in a local SELinux policy, install the **ipa-selinux** package.

Bugzilla:1461914[1]

TPM 1.2 is deprecated

The Trusted Platform Module (TPM) secure cryptoprocessor standard was updated to version 2.0 in 2016. TPM 2.0 provides many improvements over TPM 1.2, and it is not backward compatible with the previous version. TPM 1.2 is deprecated in RHEL 8, and it might be removed in the next major release.

Bugzilla:1657927[1]

crypto-policies derived properties are now deprecated

With the introduction of scopes for **crypto-policies** directives in custom policies, the following derived properties have been deprecated: **tls_cipher**, **ssh_cipher**, **ssh_group**, **ike_protocol**, and **sha1_in_dnssec**. Additionally, the use of the **protocol** property without specifying a scope is now deprecated as well. See the **crypto-policies(7)** man page for recommended replacements.

Bugzilla:2011208

10.3. SUBSCRIPTION MANAGEMENT

The --token option of the subscription-manager command is deprecated

The **--token=<TOKEN>** option of the **subscription-manager register** command is an authentication method that helps register your system to Red Hat. This option depends on capabilities offered by the entitlement server. The default entitlement server, **subscription.rhsm.redhat.com**, is planning to turn off this capability. As a consequence, attempting to use **subscription-manager register --token=<TOKEN>** might fail with the following error message:

> Token authentication not supported by the entitlement server

You can continue registering your system using other authorization methods, such as including paired options **--username** / **--password** and **--org** / **--activationkey** of the **subscription-manager register** command.

Bugzilla:2170082

10.4. SOFTWARE MANAGEMENT

rpmbuild --sign is deprecated

The **rpmbuild --sign** command is deprecated since RHEL 8.1. Using this command in future releases of Red Hat Enterprise Linux can result in an error. It is recommended that you use the **rpmsign** command instead.

Bugzilla:1688849

10.5. SHELLS AND COMMAND-LINE TOOLS

The OpenEXR component has been deprecated

The **OpenEXR** component has been deprecated. Hence, the support for the **EXR** image format has been dropped from the **imagecodecs** module.

Bugzilla:1886310

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 `hidepid=n` mount option is not supported in RHEL 8 `systemd`

The mount option **hidepid=n**, which controls who can access information in **/proc/[pid]** directories, is not compatible with **systemd** infrastructure provided in RHEL 8.

In addition, using this option might cause certain services started by **systemd** to produce SELinux AVC denial messages and prevent other operations from completing.

For more information, see the related Knowledgebase solution Is mounting /proc with "hidepid=2" recommended with RHEL7 and RHEL8?.

Bugzilla:2038929

The /usr/lib/udev/rename_device utility has been deprecated

The **udev** helper utility **/usr/lib/udev/rename_device** for renaming network interfaces has been deprecated.

Bugzilla:1875485

The ABRT tool has been deprecated

The Automatic Bug Reporting Tool (ABRT) for detecting and reporting application crashes has been deprecated in RHEL 8. As a replacement, use the **systemd-coredump** tool to log and store core dumps, which are automatically generated files after a program crashes.

Bugzilla:2055826[1]

The ReaR crontab has been deprecated

The **/etc/cron.d/rear** crontab from the **rear** package has been deprecated in RHEL 8 and will not be available in RHEL 9. The crontab checks every night whether the disk layout has changed, and runs **rear mkrescue** command if a change happened.

If you require this functionality, after an upgrade to RHEL 9, configure periodic runs of ReaR manually.

Bugzilla:2083301

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-6859

The **raw** command has been deprecated

The **raw** (**/usr/bin/raw**) command has been deprecated. Using this command in future releases of Red Hat Enterprise Linux can result in an error.

Jira:RHELPLAN-133171[1]

10.6. NETWORKING

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 8.9, the **PV_KEYv2** API is deprecated. If you use this kernel API in your application, migrate it to use the modern **netlink** API as an alternative.

Jira:RHEL-1257[1]

Network scripts are deprecated in RHEL 8

Network scripts are deprecated in Red Hat Enterprise Linux 8 and they are no longer provided by default. The basic installation provides a new version of the **ifup** and **ifdown** scripts which call the NetworkManager service through the **nmcli** tool. In Red Hat Enterprise Linux 8, to run the   **ifup** and the **ifdown** scripts, NetworkManager must be running.

Note that custom commands in **/sbin/ifup-local**, **ifdown-pre-local** and **ifdown-local** scripts are not executed.

If any of these scripts are required, the installation of the deprecated network scripts in the system is still possible with the following command:

```
# yum install network-scripts
```

The **ifup** and **ifdown** scripts link to the installed legacy network scripts.

Calling the legacy network scripts shows a warning about their deprecation.

Bugzilla:1647725[1]

The **dropwatch** tool is deprecated

The **dropwatch** tool has been deprecated. The tool will not be supported in future releases, thus it is not recommended for new deployments. As a replacement of this package, Red Hat recommends to use the **perf** command line tool.

For more information on using the **perf** command line tool, see the Getting started with Perf section on the Red Hat customer portal or the **perf** man page.

Bugzilla:1929173

The **xinetd** service has been deprecated

The **xinetd** service has been deprecated and will be removed in RHEL 9. As a replacement, use **systemd**. For further details, see How to convert xinetd service to systemd .

Bugzilla:2009113[1]

The **cgdcbxd** package is deprecated

Control group data center bridging exchange daemon (**cgdcbxd**) is a service to monitor data center bridging (DCB) netlink events and manage the **net_prio control** group subsystem. Starting with RHEL 8.5, the **cgdcbxd** package is deprecated and will be removed in the next major RHEL release.

The WEP Wi-Fi connection method is deprecated

The insecure wired equivalent privacy (WEP) Wi-Fi connection method is deprecated in RHEL 8 and will be removed in RHEL 9.0. For secure Wi-Fi connections, use the Wi-Fi Protected Access 3 (WPA3) or WPA2 connection methods.

The unsupported **xt_u32** module is now deprecated

Using the unsupported **xt_u32** module, users of **iptables** can match arbitrary 32 bits in the packet header or payload. Since RHEL 8.6, the **xt_u32** module is deprecated and will be removed in RHEL 9.

If you use **xt_u32**, migrate to the **nftables** packet filtering framework. For example, first change your firewall to use **iptables** with native matches to incrementally replace individual rules, and later use the **iptables-translate** and accompanying utilities to migrate to **nftables**. If no native match exists in **nftables**, use the raw payload matching feature of **nftables**. For details, see the **raw payload expression** section in the **nft(8)** man page.

The term **slaves** is deprecated in the **nmstate** API

Red Hat is committed to using conscious language. See details about this initiative in Making open source more inclusive. Therefore the **slaves** term is deprecated in the Nmstate API. Use the term **port** when you use **nmstatectl**.

Jira:RHELDOCS-17641

10.7. KERNEL

The **rdma_rxe** Soft-RoCE driver is deprecated

Software Remote Direct Memory Access over Converged Ethernet (Soft-RoCE), also known as RXE, is a feature that emulates Remote Direct Memory Access (RDMA). In RHEL 8, the Soft-RoCE feature is available as an unsupported Technology Preview. However, due to stability issues, this feature has been deprecated and will be removed in RHEL 9.

Bugzilla:1878207[1]

The Linux **firewire** sub-system and its associated user-space components are deprecated in RHEL 8

The **firewire** sub-system provides interfaces to use and maintain any resources on the IEEE 1394 bus. In RHEL 9, **firewire** will no longer be supported in the **kernel** package. Note that **firewire** contains several user-space components provided by the **libavc1394**, **libdc1394**, **libraw1394** packages. These packages are subject to the deprecation as well.

Bugzilla:1871863[1]

Installing RHEL for Real Time 8 using diskless boot is now deprecated

Diskless booting allows multiple systems to share a root file system through the network. While convenient, diskless boot is prone to introducing network latency in real-time workloads. With a future minor update of RHEL for Real Time 8, the diskless booting feature will no longer be supported.

Bugzilla:1748980

Kernel live patching now covers all RHEL minor releases

Since RHEL 8.1, kernel live patches have been provided for selected minor release streams of RHEL covered under the Extended Update Support (EUS) policy to remediate Critical and Important Common Vulnerabilities and Exposures (CVEs). To accommodate the maximum number of concurrently covered kernels and use cases, the support window for each live patch has been decreased from 12 to 6 months for every minor, major, and zStream version of the kernel. It means that on the day a kernel live patch is released, it will cover every minor release and scheduled errata kernel delivered in the past 6 months.

For more information about this feature, see Applying patches with kernel live patching .

For details about available kernel live patches, see Kernel Live Patch life cycles .

Bugzilla:1958250

The **crash-ptdump-command** package is deprecated

The **crash-ptdump-command** package, which is a **ptdump** extension module for the crash utility, is deprecated and might not be available in future RHEL releases. The **ptdump** command fails to retrieve the log buffer when working in the Single Range Output mode and only works in the Table of Physical Addresses (ToPA) mode. **crash-ptdump-command** is currently not maintained upstream

Bugzilla:1838927[1]

10.8. BOOT LOADER

The **kernelopts** environment variable has been deprecated

In RHEL 8, the kernel command-line parameters for systems using the GRUB bootloader were defined in the **kernelopts** environment variable. The variable was stored in the **/boot/grub2/grubenv** file for each kernel boot entry. However, storing the kernel command-line parameters using **kernelopts** was not robust. Therefore, with a future major update of RHEL, **kernelopts** will be removed and the kernel command-line parameters will be stored in the Boot Loader Specification (BLS) snippet instead.

Bugzilla:2060759

10.9. FILE SYSTEMS AND STORAGE

The **elevator** kernel command line parameter is deprecated

The **elevator** kernel command line parameter was used in earlier RHEL releases to set the disk scheduler for all devices. In RHEL 8, the parameter is deprecated.

The upstream Linux kernel has removed support for the **elevator** parameter, but it is still available in RHEL 8 for compatibility reasons.

Note that the kernel selects a default disk scheduler based on the type of device. This is typically the optimal setting. If you require a different scheduler, Red Hat recommends that you use **udev** rules or the TuneD service to configure it. Match the selected devices and switch the scheduler only for those devices.

For more information, see Setting the disk scheduler .

Bugzilla:1665295[1]

NFSv3 over UDP has been disabled

The NFS server no longer opens or listens on a User Datagram Protocol (UDP) socket by default. This change affects only NFS version 3 because version 4 requires the Transmission Control Protocol (TCP).

NFS over UDP is no longer supported in RHEL 8.

Bugzilla:1592011[1]

peripety is deprecated

The **peripety** package is deprecated since RHEL 8.3.

The Peripety storage event notification daemon parses system storage logs into structured storage events. It helps you investigate storage issues.

Bugzilla:1871953

VDO write modes other than `async` are deprecated

VDO supports several write modes in RHEL 8:

- **sync**

- **async**

- **async-unsafe**

- **auto**

Starting with RHEL 8.4, the following write modes are deprecated:

sync

Devices above the VDO layer cannot recognize if VDO is synchronous, and consequently, the devices cannot take advantage of the VDO **sync** mode.

async-unsafe

VDO added this write mode as a workaround for the reduced performance of **async** mode, which complies to Atomicity, Consistency, Isolation, and Durability (ACID). Red Hat does not recommend **async-unsafe** for most use cases and is not aware of any users who rely on it.

auto

This write mode only selects one of the other write modes. It is no longer necessary when VDO supports only a single write mode.

These write modes will be removed in a future major RHEL release.

The recommended VDO write mode is now **async**.

For more information on VDO write modes, see Selecting a VDO write mode.

Jira:RHELPLAN-70700[1]

VDO manager has been deprecated

The python-based VDO management software has been deprecated and will be removed from RHEL 9. In RHEL 9, it will be replaced by the LVM-VDO integration. Therefore, it is recommended to create VDO volumes using the **lvcreate** command.

The existing volumes created using the VDO management software can be converted using the **/usr/sbin/lvm_import_vdo** script, provided by the **lvm2** package. For more information on the LVM-VDO implementation, see Deduplicating and compressing logical volumes on RHEL .

Bugzilla:1949163

cramfs has been deprecated

Due to lack of users, the **cramfs** kernel module is deprecated. **squashfs** is recommended as an alternative solution.

Bugzilla:1794513[1]

10.10. HIGH AVAILABILITY AND CLUSTERS

pcs commands that support the **clufter** tool have been deprecated

The **pcs** commands that support the **clufter** tool for analyzing cluster configuration formats have been deprecated. These commands now print a warning that the command has been deprecated and sections related to these commands have been removed from the **pcs** help display and the **pcs(8)** man page.

The following commands have been deprecated:

- **pcs config import-cman** for importing CMAN / RHEL6 HA cluster configuration

- **pcs config export** for exporting cluster configuration to a list of **pcs** commands which recreate the same cluster

Bugzilla:1851335[1]

10.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

The **mod_php** module provided with PHP for use with the Apache HTTP Server has been deprecated

The **mod_php** module provided with PHP for use with the Apache HTTP Server in RHEL 8 is available but not enabled in the default configuration. The module is no longer available in RHEL 9.

Since RHEL 8, PHP scripts are run using the FastCGI Process Manager (**php-fpm**) by default. For more information, see Using PHP with the Apache HTTP Server .

Bugzilla:2225332

10.12. COMPILERS AND DEVELOPMENT TOOLS

The **gdb.i686** packages are deprecated

In RHEL 8.1, the 32-bit versions of the GNU Debugger (GDB), **gdb.i686**, were shipped due to a dependency problem in another package. Because RHEL 8 does not support 32-bit hardware, the **gdb.i686** packages are deprecated since RHEL 8.4. The 64-bit versions of GDB, **gdb.x86_64**, are fully capable of debugging 32-bit applications.

If you use **gdb.i686**, note the following important issues:

- The **gdb.i686** packages will no longer be updated. Users must install **gdb.x86_64** instead.

- If you have **gdb.i686** installed, installing **gdb.x86_64** will cause **yum** to report **package gdb-8.2-14.el8.x86_64 obsoletes gdb < 8.2-14.el8 provided by gdb-8.2-12.el8.i686**. This is expected. Either uninstall **gdb.i686** or pass **dnf** the **--allowerasing** option to remove **gdb.i686** and install **gdb.x8_64**.

- Users will no longer be able to install the **gdb.i686** packages on 64-bit systems, that is, those with the **libc.so.6()(64-bit)** packages.

Bugzilla:1853140[1]

libdwarf has been deprecated

The **libdwarf** library has been deprecated in RHEL 8. The library will likely not be supported in future major releases. Instead, use the **elfutils** and **libdw** libraries for applications that wish to process ELF/DWARF files.

Alternatives for the **libdwarf-tools dwarfdump** program are the **binutils readelf** program or the **elfutils eu-readelf** program, both used by passing the **--debug-dump** flag.

Bugzilla:1920624

10.13. IDENTITY MANAGEMENT

openssh-ldap has been deprecated

The **openssh-ldap** subpackage has been deprecated in Red Hat Enterprise Linux 8 and will be removed in RHEL 9. As the **openssh-ldap** subpackage is not maintained upstream, Red Hat recommends using SSSD and the **sss_ssh_authorizedkeys** helper, which integrate better with other IdM solutions and are more secure.

By default, the SSSD **ldap** and **ipa** providers read the **sshPublicKey** LDAP attribute of the user object, if available. Note that you cannot use the default SSSD configuration for the **ad** provider or IdM trusted domains to retrieve SSH public keys from Active Directory (AD), since AD does not have a default LDAP attribute to store a public key.

To allow the **sss_ssh_authorizedkeys** helper to get the key from SSSD, enable the **ssh** responder by adding **ssh** to the **services** option in the **sssd.conf** file. See the **sssd.conf(5)** man page for details.

To allow **sshd** to use **sss_ssh_authorizedkeys**, add the **AuthorizedKeysCommand /usr/bin/sss_ssh_authorizedkeys** and **AuthorizedKeysCommandUser nobody** options to the **/etc/ssh/sshd_config** file as described by the **sss_ssh_authorizedkeys(1)** man page.

Bugzilla:1871025

DES and 3DES encryption types have been removed

Due to security reasons, the Data Encryption Standard (DES) algorithm has been deprecated and disabled by default since RHEL 7. With the recent rebase of Kerberos packages, single-DES (DES) and triple-DES (3DES) encryption types have been removed from RHEL 8.

If you have configured services or users to only use DES or 3DES encryption, you might experience service interruptions such as:

- Kerberos authentication errors

- **unknown enctype** encryption errors

- Kerberos Distribution Centers (KDCs) with DES-encrypted Database Master Keys (**K/M**) fail to start

Perform the following actions to prepare for the upgrade:

1. Check if your KDC uses DES or 3DES encryption with the **krb5check** open source Python scripts. See krb5check on GitHub.

2. If you are using DES or 3DES encryption with any Kerberos principals, re-key them with a supported encryption type, such as Advanced Encryption Standard (AES). For instructions on re-keying, see Retiring DES from MIT Kerberos Documentation.

3. Test independence from DES and 3DES by temporarily setting the following Kerberos options before upgrading:

    a. In **/var/kerberos/krb5kdc/kdc.conf** on the KDC, set **supported_enctypes** and do not include **des** or **des3**.

    b. For every host, in **/etc/krb5.conf** and any files in **/etc/krb5.conf.d**, set **allow_weak_crypto** to **false**. It is false by default.

    c. For every host, in **/etc/krb5.conf** and any files in **/etc/krb5.conf.d**, set **permitted_enctypes**, **default_tgs_enctypes**, and **default_tkt_enctypes**, and do not include **des** or **des3**.

4. If you do not experience any service interruptions with the test Kerberos settings from the previous step, remove them and upgrade. You do not need those settings after upgrading to the latest Kerberos packages.

Bugzilla:1877991

The SSSD version of **libwbclient** has been removed

The SSSD implementation of the **libwbclient** package was deprecated in RHEL 8.4. As it cannot be used with recent versions of Samba, the SSSD implementation of **libwbclient** has now been removed.

Bugzilla:1947671

Standalone use of the **ctdb** service has been deprecated

Since RHEL 8.4, customers are advised to use the **ctdb** clustered Samba service only when both of the following conditions apply:

- The **ctdb** service is managed as a **pacemaker** resource with the resource-agent **ctdb**.

- The **ctdb** service uses storage volumes that contain either a GlusterFS file system provided by the Red Hat Gluster Storage product or a GFS2 file system.

The stand-alone use case of the **ctdb** service has been deprecated and will not be included in a next major release of Red Hat Enterprise Linux. For further information on support policies for Samba, see the Knowledgebase article Support Policies for RHEL Resilient Storage - ctdb General Policies .

Bugzilla:1916296[1]

Indirect AD integration with IdM via WinSync has been deprecated

WinSync is no longer actively developed in RHEL 8 due to several functional limitations:

- WinSync supports only one Active Directory (AD) domain.

- Password synchronization requires installing additional software on AD Domain Controllers.

For a more robust solution with better resource and security separation, Red Hat recommends using a **cross-forest trust** for indirect integration with Active Directory. See the Indirect integration documentation.

Jira:RHELPLAN-100400[1]

Running Samba as a PDC or BDC is deprecated

The classic domain controller mode that enabled administrators to run Samba as an NT4-like primary domain controller (PDC) and backup domain controller (BDC) is deprecated. The code and settings to configure these modes will be removed in a future Samba release.

As long as the Samba version in RHEL 8 provides the PDC and BDC modes, Red Hat supports these modes only in existing installations with Windows versions which support NT4 domains. Red Hat recommends not setting up a new Samba NT4 domain, because Microsoft operating systems later than Windows 7 and Windows Server 2008 R2 do not support NT4 domains.

If you use the PDC to authenticate only Linux users, Red Hat suggests migrating to Red Hat Identity Management (IdM) that is included in RHEL subscriptions. However, you cannot join Windows systems to an IdM domain. Note that Red Hat continues supporting the PDC functionality IdM uses in the background.

Red Hat does not support running Samba as an AD domain controller (DC).

Bugzilla:1926114

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]

Limited support for FreeRADIUS

In RHEL 8, the following external authentication modules are deprecated as part of the FreeRADIUS offering:

- The MySQL, PostgreSQL, SQlite, and unixODBC database connectors

- The **Perl** language module

- The REST API module
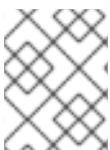
> **NOTE**
>
> The PAM authentication module and other authentication modules that are provided as part of the base package are not affected.

You can find replacements for the deprecated modules in community-supported packages, for example in the Fedora project.

In addition, the scope of support for the **freeradius** package will be limited to the following use cases in future RHEL releases:

- Using FreeRADIUS as a wireless-authentication provider with Identity Management (IdM) as the backend source of authentication. The authentication occurs through the **krb5** and LDAP authentication packages or as PAM authentication in the main FreeRADIUS package.

- Using FreeRADIUS to provide a source-of-truth for authentication in IdM, through the Python 3 authentication package.

In contrast to these deprecations, Red Hat will strengthen the support of the following external authentication modules with FreeRADIUS:

- Authentication based on **krb5** and LDAP

- **Python 3** authentication

The focus on these integration options is in close alignment with the strategic direction of Red Hat IdM.

Jira:RHELDOCS-17573

10.14. DESKTOP

The **libgnome-keyring** library has been deprecated

The **libgnome-keyring** library has been deprecated in favor of the **libsecret** library, as **libgnome-keyring** is not maintained upstream, and does not follow the necessary cryptographic policies for RHEL. The new **libsecret** library is the replacement that follows the necessary security standards.

Bugzilla:1607766[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]

10.15. GRAPHICS INFRASTRUCTURES

AGP graphics cards are no longer supported

Graphics cards using the Accelerated Graphics Port (AGP) bus are not supported in Red Hat Enterprise Linux 8. Use the graphics cards with PCI-Express bus as the recommended replacement.

Bugzilla:1569610[1]

Motif has been deprecated

The Motif widget toolkit has been deprecated in RHEL, because development in the upstream Motif community is inactive.

The following Motif packages have been deprecated, including their development and debugging variants:

- **motif**

- **openmotif**

- **openmotif21**

- **openmotif22**

Additionally, the **motif-static** package has been removed.

Red Hat recommends using the GTK toolkit as a replacement. GTK is more maintainable and provides new features compared to Motif.

Jira:RHELPLAN-98983[1]

10.16. THE WEB CONSOLE

The web console no longer supports incomplete translations

The RHEL web console no longer provides translations for languages that have translations available for less than 50 % of the Console's translatable strings. If the browser requests translation to such a language, the user interface will be in English instead.

Bugzilla:1666722

The **remotectl** command is deprecated

The **remotectl** command has been deprecated and will not be available in future releases of RHEL. You can use the **cockpit-certificate-ensure** command as a replacement. However, note that **cockpit-certificate-ensure** does not have feature parity with **remotectl**. It does not support bundled certificates and keychain files and requires them to be split out.

Jira:RHELPLAN-147538[1]

10.17. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The **geoipupdate** package has been deprecated

The **geoipupdate** package requires a third-party subscription and it also downloads proprietary content. Therefore, the **geoipupdate** package has been deprecated, and will be removed in the next major RHEL version.

Bugzilla:1874892[1]

The **network** System Role displays a deprecation warning when configuring teams on RHEL 9 nodes

The network teaming capabilities have been deprecated in RHEL 9. As a result, using the **network** RHEL System Role on an RHEL 8 control node to configure a network team on RHEL 9 nodes, shows a warning about the deprecation.

Bugzilla:2021685

Ansible Engine has been deprecated

Previous versions of RHEL 8 provided access to an Ansible Engine repository, with a limited scope of support, to enable supported RHEL Automation use cases, such as RHEL System Roles and Insights remediations. Ansible Engine has been deprecated, and Ansible Engine 2.9 will have no support after September 29, 2023. For more details on the supported use cases, see Scope of support for the Ansible Core package included in the RHEL 9 AppStream.

Users must manually migrate their systems from Ansible Engine to Ansible Core. For that, follow the steps:

Procedure

1. Check if the system is running RHEL 8.7 or a later release:

   ```
   # cat /etc/redhat-release
   ```

2. Uninstall Ansible Engine 2.9:

   ```
   # yum remove ansible
   ```

3. Disable the **ansible-2-for-rhel-8-x86_64-rpms** repository:

   ```
   # subscription-manager repos --disable
   ansible-2-for-rhel-8-x86_64-rpms
   ```

4. Install the Ansible Core package from the RHEL 8 AppStream repository:

   ```
   # yum install ansible-core
   ```

For more details, see: Using Ansible in RHEL 8.6 and later .

Bugzilla:2006081

10.18. VIRTUALIZATION

virsh iface-* commands have become deprecated

The **virsh iface-*** commands, such as **virsh iface-start** and **virsh iface-destroy**, are now deprecated, and will be removed in a future major version of RHEL. In addition, these commands frequently fail due to configuration dependencies.

Therefore, it is recommended not to use **virsh iface-*** commands for configuring and managing host network connections. Instead, use the NetworkManager program and its related management applications, such as **nmcli**.

Bugzilla:1664592[1]

virt-manager has been deprecated

The Virtual Machine Manager application, also known as **virt-manager**, has been deprecated. The RHEL web console, also known as **Cockpit**, is intended to become its replacement in a subsequent release. It is, therefore, recommended that you use the web console for managing virtualization in a GUI. Note,

however, that some features available in **virt-manager** might not be yet available in the RHEL web console.

Jira:RHELPLAN-10304[1]

Limited support for virtual machine snapshots

Creating snapshots of virtual machines (VMs) is currently only supported for VMs not using the UEFI firmware. In addition, during the snapshot operation, the QEMU monitor may become blocked, which negatively impacts the hypervisor performance for certain workloads.

Also note that the current mechanism of creating VM snapshots has been deprecated, and Red Hat does not recommend using VM snapshots in a production environment.

Bugzilla:1686057

The Cirrus VGA virtual GPU type has been deprecated

With a future major update of Red Hat Enterprise Linux, the **Cirrus VGA** GPU device will no longer be supported in KVM virtual machines. Therefore, Red Hat recommends using the **stdvga**, **virtio-vga**, or **qxl** devices instead of **Cirrus VGA**.

Bugzilla:1651994[1]

SPICE has been deprecated

The SPICE remote display protocol has become deprecated. Note that SPICE will remain supported in RHEL 8, but Red Hat recommends using alternate solutions for remote display streaming:

- For remote console access, use the VNC protocol.

- For advanced remote display functions, use third party tools such as RDP, HP RGS, or Mechdyne TGX.

Bugzilla:1849563[1]

KVM on IBM POWER has been deprecated

Using KVM virtualization on IBM POWER hardware has become deprecated. As a result, KVM on IBM POWER is still supported in RHEL 8, but will become unsupported in a future major release of RHEL.

Jira:RHELPLAN-71200[1]

SecureBoot image verification using SHA1-based signatures is deprecated

Performing SecureBoot image verification using SHA1-based signatures on UEFI (PE/COFF) executables has become deprecated. Instead, Red Hat recommends using signatures based on the SHA2 algorithm, or later.

Bugzilla:1935497[1]

Using SPICE to attach smart card readers to virtual machines has been deprecated

The SPICE remote display protocol has been deprecated in RHEL 8. Since the only recommended way to attach smart card readers to virtual machines (VMs) depends on the SPICE protocol, the usage of smart cards in VMs has also become deprecated in RHEL 8.

In a future major version of RHEL, the functionality of attaching smart card readers to VMs will only be supported by third party remote visualization solutions.

Bugzilla:2059626

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]

10.19. CONTAINERS

The Podman varlink-based API v1.0 has been removed

The Podman varlink-based API v1.0 was deprecated in a previous release of RHEL 8. Podman v2.0 introduced a new Podman v2.0 RESTful API. With the release of Podman v3.0, the varlink-based API v1.0 has been completely removed.

Jira:RHELPLAN-45858[1]

container-tools:1.0 has been deprecated

The **container-tools:1.0** module has been deprecated and will no longer receive security updates. It is recommended to use a newer supported stable module stream, such as **container-tools:2.0** or **container-tools:3.0**.

Jira:RHELPLAN-59825[1]

The **container-tools:2.0** module has been deprecated

The container-tools:2.0 module has been deprecated and will no longer receive security updates. It is recommended to use a newer supported stable module stream, such as **container-tools:3.0**.

Jira:RHELPLAN-85066[1]

Flatpak images except GIMP has been deprecated

The **rhel8/firefox-flatpak**, **rhel8/thunderbird-flatpak**, **rhel8/inkscape-flatpak**, and **rhel8/libreoffice-flatpak** RHEL 8 Flatpak Applications have been deprecated and replaced by the RHEL 9 versions. The **rhel8/gimp-flatpak** Flatpak Application is not deprecated because there is no replacement yet in RHEL 9.

Bugzilla:2142499

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-16755[1]

container-tools:3.0 has been deprecated

The **container-tools:3.0** module has been deprecated and will no longer receive security updates. To continue to build and run Linux Containers on RHEL, use a newer, stable, and supported module stream, such as **container-tools:4.0**.

For instructions on switching to a later stream, see Switching to a later stream .

Jira:RHELPLAN-146398[1]

The Inkscape and LibreOffice Flatpak images are deprecated

The **rhel9/inkscape-flatpak** and **rhel9/libreoffice-flatpak** Flatpak images, which are available as Technology Previews, have been deprecated.

Red Hat recommends the following alternatives to these images:

- To replace **rhel9/inkscape-flatpak**, use the **inkscape** RPM package.

- To replace **rhel9/libreoffice-flatpak**, see the LibreOffice deprecation release note .

Jira:RHELDOCS-17102[1]

10.20. DEPRECATED PACKAGES

This section lists packages that have been deprecated and will probably not be included in a future major release of Red Hat Enterprise Linux.

For changes to packages between RHEL 7 and RHEL 8, see Changes to packages in the *Considerations in adopting RHEL 8* document.

> **IMPORTANT**
>
> The support status of deprecated packages remains unchanged within RHEL 8. 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 8:

- 389-ds-base-legacy-tools

- abrt

- abrt-addon-ccpp

- abrt-addon-kerneloops

- abrt-addon-pstoreoops

- abrt-addon-vmcore

- abrt-addon-xorg

- abrt-cli

- abrt-console-notification

- abrt-dbus

- abrt-desktop

- abrt-gui

- abrt-gui-libs

- abrt-libs

- abrt-tui

- adobe-source-sans-pro-fonts

- adwaita-qt

- alsa-plugins-pulseaudio

- amanda

- amanda-client

- amanda-libs

- amanda-server

- ant-contrib

- antlr3

- antlr32

- aopalliance

- apache-commons-collections

- apache-commons-compress

- apache-commons-exec

- apache-commons-jxpath

- apache-commons-parent

- apache-ivy

- apache-parent

- apache-resource-bundles

- apache-sshd

- apiguardian

- aspnetcore-runtime-3.0

- aspnetcore-runtime-3.1

- aspnetcore-runtime-5.0

- aspnetcore-targeting-pack-3.0

- aspnetcore-targeting-pack-3.1

- aspnetcore-targeting-pack-5.0

- assertj-core

- authd

- auto

- autoconf213

- autogen

- autogen-libopts

- awscli

- base64coder

- batik

- batik-css

- batik-util

- bea-stax

- bea-stax-api

- bind-export-devel

- bind-export-libs

- bind-libs-lite

- bind-pkcs11

- bind-pkcs11-devel

- bind-pkcs11-libs

- bind-pkcs11-utils

- bind-sdb

- bind-sdb

- bind-sdb-chroot

- bluez-hid2hci

- boost-jam

- boost-signals

- bouncycastle

- bpg-algeti-fonts

- bpg-chveulebrivi-fonts

- bpg-classic-fonts

- bpg-courier-fonts

- bpg-courier-s-fonts

- bpg-dedaena-block-fonts

- bpg-dejavu-sans-fonts

- bpg-elite-fonts

- bpg-excelsior-caps-fonts

- bpg-excelsior-condenced-fonts

- bpg-excelsior-fonts

- bpg-fonts-common

- bpg-glaho-fonts

- bpg-gorda-fonts

- bpg-ingiri-fonts

- bpg-irubaqidze-fonts

- bpg-mikhail-stephan-fonts

- bpg-mrgvlovani-caps-fonts

- bpg-mrgvlovani-fonts

- bpg-nateli-caps-fonts

- bpg-nateli-condenced-fonts

- bpg-nateli-fonts

- bpg-nino-medium-cond-fonts

- bpg-nino-medium-fonts

- bpg-sans-fonts

- bpg-sans-medium-fonts

- bpg-sans-modern-fonts

- bpg-sans-regular-fonts

- bpg-serif-fonts

- bpg-serif-modern-fonts

- bpg-ucnobi-fonts

- brlapi-java

- bsh

- buildnumber-maven-plugin

- byaccj

- cal10n

- cbi-plugins

- cdparanoia

- cdparanoia-devel

- cdparanoia-libs

- cdrdao

- cmirror

- codehaus-parent

- codemodel

- compat-exiv2-026

- compat-guile18

- compat-hwloc1

- compat-libpthread-nonshared

- compat-libtiff3

- compat-openssl10

- compat-sap-c++-11

- compat-sap-c++-10

- compat-sap-c++-9

- createrepo_c-devel

- ctags

- ctags-etags

- custodia

- cyrus-imapd-vzic

- dbus-c++

- dbus-c++-devel

- dbus-c++-glib

- dbxtool

- dhcp-libs

- directory-maven-plugin

- directory-maven-plugin-javadoc

- dirsplit

- dleyna-connector-dbus

- dleyna-core

- dleyna-renderer

- dleyna-server

- dnssec-trigger

- dnssec-trigger-panel

- dotnet-apphost-pack-3.0

- dotnet-apphost-pack-3.1

- dotnet-apphost-pack-5.0

- dotnet-host-fxr-2.1

- dotnet-host-fxr-2.1

- dotnet-hostfxr-3.0

- dotnet-hostfxr-3.1

- dotnet-hostfxr-5.0

- dotnet-runtime-2.1

- dotnet-runtime-3.0

- dotnet-runtime-3.1

- dotnet-runtime-5.0

- dotnet-sdk-2.1

- dotnet-sdk-2.1.5xx

- dotnet-sdk-3.0

- dotnet-sdk-3.1

- dotnet-sdk-5.0

- dotnet-targeting-pack-3.0

- dotnet-targeting-pack-3.1

- dotnet-targeting-pack-5.0

- dotnet-templates-3.0

- dotnet-templates-3.1

- dotnet-templates-5.0

- dotnet5.0-build-reference-packages

- dptfxtract

- drpm

- drpm-devel

- dump

- dvd+rw-tools

- dyninst-static

- eclipse-ecf

- eclipse-ecf-core

- eclipse-ecf-runtime

- eclipse-emf

- eclipse-emf-core

- eclipse-emf-runtime

- eclipse-emf-xsd

- eclipse-equinox-osgi

- eclipse-jdt

- eclipse-license

- eclipse-p2-discovery

- eclipse-pde

- eclipse-platform

- eclipse-swt

- ed25519-java

- ee4j-parent

- elfutils-devel-static

- elfutils-libelf-devel-static

- enca

- enca-devel

- environment-modules-compat

- evince-browser-plugin

- exec-maven-plugin

- farstream02

- felix-gogo-command

- felix-gogo-runtime

- felix-gogo-shell

- felix-scr

- felix-osgi-compendium

- felix-osgi-core

- felix-osgi-foundation

- felix-parent

- file-roller

- fipscheck

- fipscheck-devel

- fipscheck-lib

- firewire

- fonts-tweak-tool

- forge-parent

- freeradius-mysql

- freeradius-perl

- freeradius-postgresql

- freeradius-rest

- freeradius-sqlite

- freeradius-unixODBC

- fuse-sshfs

- fusesource-pom

- future

- gamin

- gamin-devel

- gavl

- gcc-toolset-9

- gcc-toolset-9-annobin

- gcc-toolset-9-build

- gcc-toolset-9-perftools

- gcc-toolset-9-runtime

- gcc-toolset-9-toolchain

- gcc-toolset-10

- gcc-toolset-10-annobin

- gcc-toolset-10-binutils

- gcc-toolset-10-binutils-devel

- gcc-toolset-10-build

- gcc-toolset-10-dwz

- gcc-toolset-10-dyninst

- gcc-toolset-10-dyninst-devel

- gcc-toolset-10-elfutils

- gcc-toolset-10-elfutils-debuginfod-client

- gcc-toolset-10-elfutils-debuginfod-client-devel

- gcc-toolset-10-elfutils-devel

- gcc-toolset-10-elfutils-libelf

- gcc-toolset-10-elfutils-libelf-devel

- gcc-toolset-10-elfutils-libs

- gcc-toolset-10-gcc

- gcc-toolset-10-gcc-c++

- gcc-toolset-10-gcc-gdb-plugin

- gcc-toolset-10-gcc-gfortran

- gcc-toolset-10-gdb

- gcc-toolset-10-gdb-doc

- gcc-toolset-10-gdb-gdbserver

- gcc-toolset-10-libasan-devel

- gcc-toolset-10-libatomic-devel

- gcc-toolset-10-libitm-devel

- gcc-toolset-10-liblsan-devel

- gcc-toolset-10-libquadmath-devel

- gcc-toolset-10-libstdc++-devel

- gcc-toolset-10-libstdc++-docs

- gcc-toolset-10-libtsan-devel

- gcc-toolset-10-libubsan-devel

- gcc-toolset-10-ltrace

- gcc-toolset-10-make

- gcc-toolset-10-make-devel

- gcc-toolset-10-perftools

- gcc-toolset-10-runtime

- gcc-toolset-10-strace

- gcc-toolset-10-systemtap

- gcc-toolset-10-systemtap-client

- gcc-toolset-10-systemtap-devel

- gcc-toolset-10-systemtap-initscript

- gcc-toolset-10-systemtap-runtime

- gcc-toolset-10-systemtap-sdt-devel

- gcc-toolset-10-systemtap-server

- gcc-toolset-10-toolchain

- gcc-toolset-10-valgrind

- gcc-toolset-10-valgrind-devel

- gcc-toolset-11-make-devel

- GConf2

- GConf2-devel

- gegl

- genisoimage

- genwqe-tools

- genwqe-vpd

- genwqe-zlib

- genwqe-zlib-devel

- geoipupdate

- geronimo-annotation

- geronimo-jms

- geronimo-jpa

- geronimo-parent-poms

- gfbgraph

- gflags

- gflags-devel

- glassfish-annotation-api

- glassfish-el

- glassfish-fastinfoset

- glassfish-jaxb-core

- glassfish-jaxb-txw2

- glassfish-jsp

- glassfish-jsp-api

- glassfish-legal

- glassfish-master-pom

- glassfish-servlet-api

- glew-devel

- glib2-fam

- glog

- glog-devel

- gmock

- gmock-devel

- gnome-abrt

- gnome-boxes

- gnome-menus-devel

- gnome-online-miners

- gnome-shell-extension-disable-screenshield

- gnome-shell-extension-horizontal-workspaces

- gnome-shell-extension-no-hot-corner

- gnome-shell-extension-window-grouper

- gnome-themes-standard

- gnu-free-fonts-common

- gnu-free-mono-fonts

- gnu-free-sans-fonts

- gnu-free-serif-fonts

- gnupg2-smime

- gnuplot

- gnuplot-common

- gobject-introspection-devel

- google-gson

- google-noto-sans-syriac-eastern-fonts

- google-noto-sans-syriac-estrangela-fonts

- google-noto-sans-syriac-western-fonts

- google-noto-sans-tibetan-fonts

- google-noto-sans-ui-fonts

- gphoto2

- gsl-devel

- gssntlmssp

- gtest

- gtest-devel

- gtkmm24

- gtkmm24-devel

- gtkmm24-docs

- gtksourceview3

- gtksourceview3-devel

- gtkspell

- gtkspell-devel

- gtkspell3

- guile

- gutenprint-gimp

- gutenprint-libs-ui

- gvfs-afc

- gvfs-afp

- gvfs-archive

- hamcrest-core

- hawtjni

- hawtjni

- hawtjni-runtime

- HdrHistogram

- HdrHistogram-javadoc

- highlight-gui

- hivex-devel

- hostname

- hplip-gui

- httpcomponents-project

- hwloc-plugins

- hyphen-fo

- hyphen-grc

- hyphen-hsb

- hyphen-ia

- hyphen-is

- hyphen-ku

- hyphen-mi

- hyphen-mn

- hyphen-sa

- hyphen-tk

- ibus-sayura

- icedax

- icu4j

- idm-console-framework

- inkscape

- inkscape-docs

- inkscape-view

- iptables

- ipython

- isl

- isl-devel

- isorelax

- istack-commons-runtime

- istack-commons-tools

- iwl3945-firmware

- iwl4965-firmware

- iwl6000-firmware

- jacoco

- jaf

- jaf-javadoc

- jakarta-oro

- janino

- jansi-native

- jarjar

- java-1.8.0-ibm

- java-1.8.0-ibm-demo

- java-1.8.0-ibm-devel

- java-1.8.0-ibm-headless

- java-1.8.0-ibm-jdbc

- java-1.8.0-ibm-plugin

- java-1.8.0-ibm-src

- java-1.8.0-ibm-webstart

- java-1.8.0-openjdk-accessibility

- java-1.8.0-openjdk-accessibility-slowdebug

- java_cup

- java-atk-wrapper

- javacc

- javacc-maven-plugin

- javaewah

- javaparser

- javapoet

- javassist

- javassist-javadoc

- jaxen

- jboss-annotations-1.2-api

- jboss-interceptors-1.2-api

- jboss-logmanager

- jboss-parent

- jctools

- jdepend

- jdependency

- jdom

- jdom2

- jetty

- jetty-continuation

- jetty-http

- jetty-io

- jetty-security

- jetty-server

- jetty-servlet

- jetty-util

- jffi

- jflex

- jgit

- jline

- jmc

- jnr-netdb

- jolokia-jvm-agent

- js-uglify

- jsch

- json_simple

- jss-javadoc

- jtidy

- junit5

- jvnet-parent

- jzlib

- kernel-cross-headers

- ksc

- kurdit-unikurd-web-fonts

- kyotocabinet-libs

- ldapjdk-javadoc

- lensfun

- lensfun-devel

- lftp-scripts

- libaec

- libaec-devel

- libappindicator-gtk3

- libappindicator-gtk3-devel

- libatomic-static

- libavc1394

- libblocksruntime

- libcacard

- libcacard-devel

- libcgroup

- libcgroup-tools

- libchamplain

- libchamplain-devel

- libchamplain-gtk

- libcroco

- libcroco-devel

- libcxl

- libcxl-devel

- libdap

- libdap-devel

- libdazzle-devel

- libdbusmenu

- libdbusmenu-devel

- libdbusmenu-doc

- libdbusmenu-gtk3

- libdbusmenu-gtk3-devel

- libdc1394

- libdnet

- libdnet-devel

- libdv

- libdwarf

- libdwarf-devel

- libdwarf-static

- libdwarf-tools

- libeasyfc

- libeasyfc-gobject

- libepubgen-devel

- libertas-sd8686-firmware

- libertas-usb8388-firmware

- libertas-usb8388-olpc-firmware

- libgdither

- libGLEW

- libgovirt

- libguestfs-benchmarking

- libguestfs-devel

- libguestfs-gfs2

- libguestfs-gobject

- libguestfs-gobject-devel

- libguestfs-java

- libguestfs-java-devel

- libguestfs-javadoc

- libguestfs-man-pages-ja

- libguestfs-man-pages-uk

- libguestfs-tools

- libguestfs-tools-c

- libhugetlbfs

- libhugetlbfs-devel

- libhugetlbfs-utils

- libIDL

- libIDL-devel

- libidn

- libiec61883

- libindicator-gtk3

- libindicator-gtk3-devel

- libiscsi-devel

- libjose-devel

- libkkc

- libkkc-common

- libkkc-data

- libldb-devel

- liblogging

- libluksmeta-devel

- libmalaga

- libmcpp

- libmemcached

- libmemcached-libs

- libmetalink

- libmodulemd1

- libmongocrypt

- libmtp-devel

- libmusicbrainz5

- libmusicbrainz5-devel

- libnbd-devel

- liboauth

- liboauth-devel

- libpfm-static

- libpng12

- libpurple

- libpurple-devel

- libraw1394

- libreport-plugin-mailx

- libreport-plugin-rhtsupport

- libreport-plugin-ureport

- libreport-rhel

- libreport-rhel-bugzilla

- librpmem

- librpmem-debug

- librpmem-devel

- libsass

- libsass-devel

- libselinux-python

- libsqlite3x

- libtalloc-devel

- libtar

- libtdb-devel

- libtevent-devel

- libtpms-devel

- libunwind

- libusal

- libvarlink

- libverto-libevent

- libvirt-admin

- libvirt-bash-completion

- libvirt-daemon-driver-storage-gluster

- libvirt-daemon-driver-storage-iscsi-direct

- libvirt-devel

- libvirt-docs

- libvirt-gconfig

- libvirt-gobject

- libvirt-lock-sanlock

- libvirt-wireshark

- libvmem

- libvmem-debug

- libvmem-devel

- libvmmalloc

- libvmmalloc-debug

- libvmmalloc-devel

- libvncserver

- libwinpr-devel

- libwmf

- libwmf-devel

- libwmf-lite

- libXNVCtrl

- libyami

- log4j12

- log4j12-javadoc

- lohit-malayalam-fonts

- lohit-nepali-fonts

- lorax-composer

- lua-guestfs

- lucene

- lucene-analysis

- lucene-analyzers-smartcn

- lucene-queries

- lucene-queryparser

- lucene-sandbox

- lz4-java

- lz4-java-javadoc

- mailman

- mailx

- make-devel

- malaga

- malaga-suomi-voikko

- marisa

- maven-antrun-plugin

- maven-assembly-plugin

- maven-clean-plugin

- maven-dependency-analyzer

- maven-dependency-plugin

- maven-doxia

- maven-doxia-sitetools

- maven-install-plugin

- maven-invoker

- maven-invoker-plugin

- maven-parent

- maven-plugins-pom

- maven-reporting-api

- maven-reporting-impl

- maven-resolver-api

- maven-resolver-connector-basic

- maven-resolver-impl

- maven-resolver-spi

- maven-resolver-transport-wagon

- maven-resolver-util

- maven-scm

- maven-script-interpreter

- maven-shade-plugin

- maven-shared

- maven-verifier

- maven-wagon-file

- maven-wagon-http

- maven-wagon-http-shared

- maven-wagon-provider-api

- maven2

- meanwhile

- mercurial

- mercurial-hgk

- metis

- metis-devel

- mingw32-bzip2

- mingw32-bzip2-static

- mingw32-cairo

- mingw32-expat

- mingw32-fontconfig

- mingw32-freetype

- mingw32-freetype-static

- mingw32-gstreamer1

- mingw32-harfbuzz

- mingw32-harfbuzz-static

- mingw32-icu

- mingw32-libjpeg-turbo

- mingw32-libjpeg-turbo-static

- mingw32-libpng

- mingw32-libpng-static

- mingw32-libtiff

- mingw32-libtiff-static

- mingw32-openssl

- mingw32-readline

- mingw32-sqlite

- mingw32-sqlite-static

- mingw64-adwaita-icon-theme

- mingw64-bzip2

- mingw64-bzip2-static

- mingw64-cairo

- mingw64-expat

- mingw64-fontconfig

- mingw64-freetype

- mingw64-freetype-static

- mingw64-gstreamer1

- mingw64-harfbuzz

- mingw64-harfbuzz-static

- mingw64-icu

- mingw64-libjpeg-turbo

- mingw64-libjpeg-turbo-static

- mingw64-libpng

- mingw64-libpng-static

- mingw64-libtiff

- mingw64-libtiff-static

- mingw64-nettle

- mingw64-openssl

- mingw64-readline

- mingw64-sqlite

- mingw64-sqlite-static

- modello

- mojo-parent

- mongo-c-driver

- mousetweaks

- mozjs52

- mozjs52-devel

- mozjs60

- mozjs60-devel

- mozvoikko

- msv-javadoc

- msv-manual

- munge-maven-plugin

- mythes-mi

- mythes-ne

- nafees-web-naskh-fonts

- nbd

- nbdkit-devel

- nbdkit-example-plugins

- nbdkit-gzip-plugin

- nbdkit-plugin-python-common

- nbdkit-plugin-vddk

- ncompress

- ncurses-compat-libs

- net-tools

- netcf

- netcf-devel

- netcf-libs

- network-scripts

- network-scripts-ppp

- nkf

- nodejs-devel

- nodejs-packaging

- nss_nis

- nss-pam-ldapd

- objectweb-asm

- objectweb-asm-javadoc

- objectweb-pom

- ocaml-bisect-ppx

- ocaml-camlp4

- ocaml-camlp4-devel

- ocaml-lwt

- ocaml-mmap

- ocaml-ocplib-endian

- ocaml-ounit

- ocaml-result

- ocaml-seq

- opencryptoki-tpmtok

- opencv-contrib

- opencv-core

- opencv-devel

- openhpi

- openhpi-libs

- OpenIPMI-perl

- openssh-cavs

- openssh-ldap

- openssl-ibmpkcs11

- opentest4j

- os-maven-plugin

- pakchois

- pandoc

- paps-libs

- paranamer

- parfait

- parfait-examples

- parfait-javadoc

- pcp-parfait-agent

- pcp-pmda-rpm

- pcp-pmda-vmware

- pcsc-lite-doc

- peripety

- perl-B-Debug

- perl-B-Lint

- perl-Class-Factory-Util

- perl-Class-ISA

- perl-DateTime-Format-HTTP

- perl-DateTime-Format-Mail

- perl-File-CheckTree

- perl-homedir

- perl-libxml-perl

- perl-Locale-Codes

- perl-Mozilla-LDAP

- perl-NKF

- perl-Object-HashBase-tools

- perl-Package-DeprecationManager

- perl-Pod-LaTeX

- perl-Pod-Plainer

- perl-prefork

- perl-String-CRC32

- perl-SUPER

- perl-Sys-Virt

- perl-tests

- perl-YAML-Syck

- phodav

- php-recode

- php-xmlrpc

- pidgin

- pidgin-devel

- pidgin-sipe

- pinentry-emacs

- pinentry-gtk

- pipewire0.2-devel

- pipewire0.2-libs

- platform-python-coverage

- plexus-ant-factory

- plexus-bsh-factory

- plexus-cli

- plexus-component-api

- plexus-component-factories-pom

- plexus-components-pom

- plexus-i18n

- plexus-interactivity

- plexus-pom

- plexus-velocity

- plymouth-plugin-throbgress

- pmreorder

- postgresql-test-rpm-macros

- powermock

- prometheus-jmx-exporter

- prometheus-jmx-exporter-openjdk11

- ptscotch-mpich

- ptscotch-mpich-devel

- ptscotch-mpich-devel-parmetis

- ptscotch-openmpi

- ptscotch-openmpi-devel

- purple-sipe

- pygobject2-doc

- pygtk2

- pygtk2-codegen

- pygtk2-devel

- pygtk2-doc

- python-nose-docs

- python-nss-doc

- python-podman-api

- python-psycopg2-doc

- python-pymongo-doc

- python-redis

- python-schedutils

- python-slip

- python-sqlalchemy-doc

- python-varlink

- python-virtualenv-doc

- python2-backports

- python2-backports-ssl_match_hostname

- python2-bson

- python2-coverage

- python2-docs

- python2-docs-info

- python2-funcsigs

- python2-ipaddress

- python2-mock

- python2-nose

- python2-numpy-doc

- python2-psycopg2-debug

- python2-psycopg2-tests

- python2-pymongo

- python2-pymongo-gridfs

- python2-pytest-mock

- python2-sqlalchemy

- python2-tools

- python2-virtualenv

- python3-bson

- python3-click

- python3-coverage

- python3-cpio

- python3-custodia

- python3-docs

- python3-flask

- python3-gevent

- python3-gobject-base

- python3-hivex

- python3-html5lib

- python3-hypothesis

- python3-ipatests

- python3-itsdangerous

- python3-jwt

- python3-libguestfs

- python3-mock

- python3-networkx-core

- python3-nose

- python3-nss

- python3-openipmi

- python3-pillow

- python3-ptyprocess

- python3-pydbus

- python3-pymongo

- python3-pymongo-gridfs

- python3-pyOpenSSL

- python3-pytoml

- python3-reportlab

- python3-schedutils

- python3-scons

- python3-semantic_version

- python3-slip

- python3-slip-dbus

- python3-sqlalchemy

- python3-syspurpose

- python3-virtualenv

- python3-webencodings

- python3-werkzeug

- python38-asn1crypto

- python38-numpy-doc

- python38-psycopg2-doc

- python38-psycopg2-tests

- python39-numpy-doc

- python39-psycopg2-doc

- python39-psycopg2-tests

- qemu-kvm-block-gluster

- qemu-kvm-block-iscsi

- qemu-kvm-block-ssh

- qemu-kvm-hw-usbredir

- qemu-kvm-device-display-virtio-gpu-gl

- qemu-kvm-device-display-virtio-gpu-pci-gl

- qemu-kvm-device-display-virtio-vga-gl

- qemu-kvm-tests

- qpdf

- qpdf-doc

- qpid-proton

- qrencode

- qrencode-devel

- qrencode-libs

- qt5-qtcanvas3d

- qt5-qtcanvas3d-examples

- rarian

- rarian-compat

- re2c

- recode

- redhat-lsb

- redhat-lsb-core

- redhat-lsb-cxx

- redhat-lsb-desktop

- redhat-lsb-languages

- redhat-lsb-printing

- redhat-lsb-submod-multimedia

- redhat-lsb-submod-security

- redhat-lsb-supplemental

- redhat-lsb-trialuse

- redhat-menus

- redhat-support-lib-python

- redhat-support-tool

- reflections

- regexp

- relaxngDatatype

- rhsm-gtk

- rpm-plugin-prioreset

- rpmemd

- rsyslog-udpspoof

- ruby-hivex

- ruby-libguestfs

- rubygem-abrt

- rubygem-abrt-doc

- rubygem-bson

- rubygem-bson-doc

- rubygem-bundler-doc

- rubygem-mongo

- rubygem-mongo-doc

- rubygem-net-telnet

- rubygem-xmlrpc

- s390utils-cmsfs

- samba-pidl

- samba-test

- samba-test-libs

- samyak-devanagari-fonts

- samyak-fonts-common

- samyak-gujarati-fonts

- samyak-malayalam-fonts

- samyak-odia-fonts

- samyak-tamil-fonts

- sane-frontends

- sanlk-reset

- sat4j

- scala

- scotch

- scotch-devel

- SDL_sound

- selinux-policy-minimum

- sendmail

- sgabios

- sgabios-bin

- shrinkwrap

- sisu-inject

- sisu-mojos

- sisu-plexus

- skkdic

- SLOF

- smc-anjalioldlipi-fonts

- smc-dyuthi-fonts

- smc-fonts-common

- smc-kalyani-fonts

- smc-raghumalayalam-fonts

- smc-suruma-fonts

- softhsm-devel

- sonatype-oss-parent

- sonatype-plugins-parent

- sos-collector

- sparsehash-devel

- spax

- spec-version-maven-plugin

- spice

- spice-client-win-x64

- spice-client-win-x86

- spice-glib

- spice-glib-devel

- spice-gtk

- spice-gtk-tools

- spice-gtk3

- spice-gtk3-devel

- spice-gtk3-vala

- spice-parent

- spice-protocol

- spice-qxl-wddm-dod

- spice-server

- spice-server-devel

- spice-qxl-xddm

- spice-server

- spice-streaming-agent

- spice-vdagent-win-x64

- spice-vdagent-win-x86

- sssd-libwbclient

- star

- stax-ex

- stax2-api

- stringtemplate

- stringtemplate4

- subscription-manager-initial-setup-addon

- subscription-manager-migration

- subscription-manager-migration-data

- subversion-javahl

- SuperLU

- SuperLU-devel

- supermin-devel

- swig

- swig-doc

- swig-gdb

- swtpm-devel

- swtpm-tools-pkcs11
- system-storage-manager
- tcl-brlapi
- testng
- tibetan-machine-uni-fonts
- timedatex
- tpm-quote-tools
- tpm-tools
- tpm-tools-pkcs11
- treelayout
- trousers
- trousers-lib
- tuned-profiles-compat
- tuned-profiles-nfv-host-bin
- tuned-utils-systemtap
- tycho
- uglify-js
- unbound-devel
- univocity-output-tester
- univocity-parsers
- usbguard-notifier
- usbredir-devel
- utf8cpp
- uthash
- velocity
- vinagre
- vino
- virt-dib
- virt-p2v-maker

- vm-dump-metrics-devel

- weld-parent

- wodim

- woodstox-core

- wqy-microhei-fonts

- wqy-unibit-fonts

- xdelta

- xmlgraphics-commons

- xmlstreambuffer

- xinetd

- xorg-x11-apps

- xorg-x11-drv-qxl

- xorg-x11-server-Xspice

- xpp3

- xsane-gimp

- xsom

- xz-java

- xz-java-javadoc

- yajl-devel

- yp-tools

- ypbind

- ypserv

10.21. DEPRECATED AND UNMAINTAINED DEVICES

This section lists devices (drivers, adapters) that

- continue to be supported until the end of life of RHEL 8 but will likely not be supported in future major releases of this product and are not recommended for new deployments. Support for devices other than those listed remains unchanged. These are **deprecated** devices.

- are available but are no longer being tested or updated on a routine basis in RHEL 8. Red Hat may fix serious bugs, including security bugs, at its discretion. These devices should no longer be used in production, and it is likely they will be disabled in the next major release. These are **unmaintained** devices.

PCI device IDs are in the format of *vendor:device:subvendor:subdevice*. If no device ID is listed, all devices associated with the corresponding driver have been deprecated. To check the PCI IDs of the hardware on your system, run the **lspci -nn** command.

Table 10.1. Deprecated devices

Device ID	Driver	Device name
	bnx2	QLogic BCM5706/5708/5709/5716 Driver
	hpsa	Hewlett-Packard Company: Smart Array Controllers
0x10df:0x0724	lpfc	Emulex Corporation: OneConnect FCoE Initiator (Skyhawk)
0x10df:0xe200	lpfc	Emulex Corporation: LPe15000/LPe16000 Series 8Gb/16Gb Fibre Channel Adapter
0x10df:0xf011	lpfc	Emulex Corporation: Saturn: LightPulse Fibre Channel Host Adapter
0x10df:0xf015	lpfc	Emulex Corporation: Saturn: LightPulse Fibre Channel Host Adapter
0x10df:0xf100	lpfc	Emulex Corporation: LPe12000 Series 8Gb Fibre Channel Adapter
0x10df:0xfc40	lpfc	Emulex Corporation: Saturn-X: LightPulse Fibre Channel Host Adapter
0x10df:0xe220	be2net	Emulex Corporation: OneConnect NIC (Lancer)
0x1000:0x005b	megaraid_sas	Broadcom / LSI: MegaRAID SAS 2208 [Thunderbolt]
0x1000:0x006E	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2
0x1000:0x0080	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0081	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0082	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0083	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0084	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0085	mpt3sas	Broadcom / LSI: SAS2208 PCI-Express Fusion-MPT SAS-2
0x1000:0x0086	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2

Device ID	Driver	Device name
0x1000:0x0087	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2
	myri10ge	Myricom 10G driver (10GbE)
	netxen_nic	QLogic/NetXen (1/10) GbE Intelligent Ethernet Driver
0x1077:0x2031	qla2xxx	QLogic Corp.: ISP8324-based 16Gb Fibre Channel to PCI Express Adapter
0x1077:0x2532	qla2xxx	QLogic Corp.: ISP2532-based 8Gb Fibre Channel to PCI Express HBA
0x1077:0x8031	qla2xxx	QLogic Corp.: 8300 Series 10GbE Converged Network Adapter (FCoE)
	qla3xxx	QLogic ISP3XXX Network Driver v2.03.00-k5
0x1924:0x0803	sfc	Solarflare Communications: SFC9020 10G Ethernet Controller
0x1924:0x0813	sfc	Solarflare Communications: SFL9021 10GBASE-T Ethernet Controller
	Soft-RoCE (rdma_rxe)	
	HNS-RoCE	HNS GE/10GE/25GE/50GE/100GE RDMA Network Controller
	liquidio	Cavium LiquidIO Intelligent Server Adapter Driver
	liquidio_vf	Cavium LiquidIO Intelligent Server Adapter Virtual Function Driver

Table 10.2. Unmaintained devices

Device ID	Driver	Device name
	e1000	Intel® PRO/1000 Network Driver
	mptbase	Fusion MPT SAS Host driver

Device ID	Driver	Device name
	mptsas	Fusion MPT SAS Host driver
	mptscsih	Fusion MPT SCSI Host driver
	mptspi	Fusion MPT SAS Host driver
0x1000:0x0071 [a]	megaraid_sas	Broadcom / LSI: MR SAS HBA 2004
0x1000:0x0073 [a]	megaraid_sas	Broadcom / LSI: MegaRAID SAS 2008 [Falcon]
0x1000:0x0079 [a]	megaraid_sas	Broadcom / LSI: MegaRAID SAS 2108 [Liberator]
	nvmet_tcp	NVMe/TCP target driver
	nvmet-fc	NVMe/Fabrics FC target driver

[a] Disabled in RHEL 8.0, re-enabled in RHEL 8.4 due to customer requests.

CHAPTER 11. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 8.9.

11.1. INSTALLER AND IMAGE CREATION

Installation fails on IBM Power 10 systems with LPAR and secure boot enabled

RHEL installer is not integrated with static key secure boot on IBM Power 10 systems. Consequently, when logical partition (LPAR) is enabled with the secure boot option, the installation fails with the error, **Unable to proceed with RHEL-x.x Installation**.

To work around this problem, install RHEL without enabling secure boot. After booting the system:

1. Copy the signed Kernel into the PReP partition using the **dd** command.

2. Restart the system and enable secure boot.

Once the firmware verifies the bootloader and the kernel, the system boots up successfully.

For more information, see https://www.ibm.com/support/pages/node/6528884

Bugzilla:2025814[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

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]

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

Network access is not enabled by default in the installation program

Several installation features require network access, for example, registration of a system using the Content Delivery Network (CDN), NTP server support, and network installation sources. However, network access is not enabled by default, and as a result, these features cannot be used until network access is enabled.

To work around this problem, add **ip=dhcp** to boot options to enable network access when the installation starts. Optionally, passing a Kickstart file or a repository located on the network using boot options also resolves the problem. As a result, the network-based installation features can be used.

Bugzilla:1757877[1]

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

IBM Power systems with **HASH MMU** mode fail to boot with memory allocation failures

IBM Power Systems with **HASH memory allocation unit (MMU)** mode support **kdump** up to a maximum of 192 cores. Consequently, the system fails to boot with memory allocation failures if **kdump** is enabled on more than 192 cores. This limitation is due to RMA memory allocations during early boot in **HASH MMU** mode. To work around this problem, use the **Radix MMU** mode with **fadump** enabled instead of using **kdump**.

Bugzilla:2028361[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-4744

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

11.2. SECURITY

sshd -T provides inaccurate information about Ciphers, MACs and KeX algorithms

The output of the **sshd -T** command does not contain the system-wide crypto policy configuration or other options that could come from an environment file in **/etc/sysconfig/sshd** and that are applied as arguments on the **sshd** command. This occurs because the upstream OpenSSH project did not support the Include directive to support Red-Hat-provided cryptographic defaults in RHEL 8. Crypto policies are applied as command-line arguments to the **sshd** executable in the **sshd.service** unit during the service's start by using an **EnvironmentFile**. To work around the problem, use the **source** command with the environment file and pass the crypto policy as an argument to the **sshd** command, as in **sshd -T $CRYPTO_POLICY**. For additional information, see Ciphers, MACs or KeX algorithms differ from sshd -T to what is provided by current crypto policy level . As a result, the output from **sshd -T** matches the currently configured crypto policy.

Bugzilla:2044354[1]

RHV hypervisor may not work correctly when hardening the system during installation

When installing Red Hat Virtualization Hypervisor (RHV-H) and applying the Red Hat Enterprise Linux 8 STIG profile, OSCAP Anaconda Add-on may harden the system as RHEL instead of RVH-H and remove essential packages for RHV-H. Consequently, the RHV hypervisor may not work. To work around the problem, install the RHV-H system without applying any profile hardening, and after the installation is complete, apply the profile by using OpenSCAP. As a result, the RHV hypervisor works correctly.

Jira:RHEL-1826

CVE OVAL feeds are now only in the compressed format, and data streams are not in the SCAP 1.3 standard

Red Hat provides CVE OVAL feeds in the bzip2-compressed format and are no longer available in the XML file format. Because referencing compressed content is not standardized in the Security Content Automation Protocol (SCAP) 1.3 specification, third-party SCAP scanners can have problems scanning rules that use the feed.

Bugzilla:2028428

Certain Rsyslog priority strings do not work correctly

Support for the GnuTLS priority string for **imtcp** that allows fine-grained control over encryption is not complete. Consequently, the following priority strings do not work properly in the Rsyslog remote logging application:

```
NONE:+VERS-ALL:-VERS-TLS1.3:+MAC-ALL:+DHE-RSA:+AES-256-GCM:+SIGN-RSA-
SHA384:+COMP-ALL:+GROUP-ALL
```

To work around this problem, use only correctly working priority strings:

```
NONE:+VERS-ALL:-VERS-TLS1.3:+MAC-ALL:+ECDHE-RSA:+AES-128-CBC:+SIGN-RSA-
SHA1:+COMP-ALL:+GROUP-ALL
```

As a result, current configurations must be limited to the strings that work correctly.

Bugzilla:1679512

Server with GUI and Workstation installations are not possible with CIS Server profiles

The CIS Server Level 1 and Level 2 security profiles are not compatible with the **Server with GUI** and **Workstation** software selections. As a consequence, a RHEL 8 installation with the **Server with GUI** software selection and CIS Server profiles is not possible. An attempted installation using the CIS Server Level 1 or Level 2 profiles and either of these software selections will generate the error message:

```
package xorg-x11-server-common has been added to the list of excluded packages, but it can't be
removed from the current software selection without breaking the installation.
```

If you need to align systems with the **Server with GUI** or **Workstation** software selections according to CIS benchmarks, use the CIS Workstation Level 1 or Level 2 profiles instead.

Bugzilla:1843932

Kickstart uses org_fedora_oscap instead of com_redhat_oscap in RHEL 8

The Kickstart references the Open Security Content Automation Protocol (OSCAP) Anaconda add-on as **org_fedora_oscap** instead of **com_redhat_oscap**, which might cause confusion. This is necessary to keep compatibility with Red Hat Enterprise Linux 7.

Bugzilla:1665082[1]

libvirt overrides xccdf_org.ssgproject.content_rule_sysctl_net_ipv4_conf_all_forwarding

The **libvirt** virtualization framework enables IPv4 forwarding whenever a virtual network with a forward mode of **route** or **nat** is started. This overrides the configuration by the **xccdf_org.ssgproject.content_rule_sysctl_net_ipv4_conf_all_forwarding** rule, and subsequent compliance scans report the **fail** result when assessing this rule.

Apply one of these scenarios to work around the problem:

- Uninstall the **libvirt** packages if your scenario does not require them.

- Change the forwarding mode of virtual networks created by **libvirt**.

- Remove the **xccdf_org.ssgproject.content_rule_sysctl_net_ipv4_conf_all_forwarding** rule by tailoring your profile.

Bugzilla:2118758

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]

OpenSSL in FIPS mode accepts only specific D-H parameters

In FIPS mode, TLS clients that use OpenSSL return a **bad dh value** error and abort TLS connections to servers that use manually generated parameters. This is because OpenSSL, when configured to work in compliance with FIPS 140-2, works only with Diffie-Hellman parameters compliant to NIST SP 800-56A rev3 Appendix D (groups 14, 15, 16, 17, and 18 defined in RFC 3526 and with groups defined in RFC 7919). Also, servers that use OpenSSL ignore all other parameters and instead select known parameters of similar size. To work around this problem, use only the compliant groups.

Bugzilla:1810911[1]

crypto-policies incorrectly allow Camellia ciphers

The RHEL 8 system-wide cryptographic policies should disable Camellia ciphers in all policy levels, as stated in the product documentation. However, the Kerberos protocol enables the ciphers by default.

To work around the problem, apply the **NO-CAMELLIA** subpolicy:

```
# update-crypto-policies --set DEFAULT:NO-CAMELLIA
```

In the previous command, replace **DEFAULT** with the cryptographic level name if you have switched from **DEFAULT** previously.

As a result, Camellia ciphers are correctly disallowed across all applications that use system-wide crypto policies only when you disable them through the workaround.

Bugzilla:1919155

OpenSC might not detect CardOS V5.3 card objects correctly

The OpenSC toolkit does not correctly read cache from different PKCS #15 file offsets used in some CardOS V5.3 cards. Consequently, OpenSC might not be able to list card objects and prevent using them from different applications.

To work around the problem, turn off file caching by setting the **use_file_caching = false** option in the **/etc/opensc.conf** file.

Jira:RHEL-4077

Smart-card provisioning process through OpenSC **pkcs15-init** does not work properly

The **file_caching** option is enabled in the default OpenSC configuration, and the file caching functionality does not handle some commands from the **pkcs15-init** tool properly. Consequently, the smart-card provisioning process through OpenSC fails.

To work around the problem, add the following snippet to the **/etc/opensc.conf** file:

```
app pkcs15-init {
    framework pkcs15 {
        use_file_caching = false;
    }
}
```

The smart-card provisioning through **pkcs15-init** only works if you apply the previously described workaround.

Bugzilla:1947025

Connections to servers with SHA-1 signatures do not work with GnuTLS

SHA-1 signatures in certificates are rejected by the GnuTLS secure communications library as insecure. Consequently, applications that use GnuTLS as a TLS backend cannot establish a TLS connection to peers that offer such certificates. This behavior is inconsistent with other system cryptographic libraries.

To work around this problem, upgrade the server to use certificates signed with SHA-256 or stronger hash, or switch to the LEGACY policy.

Bugzilla:1628553[1]

libselinux-python is available only through its module

The **libselinux-python** package contains only Python 2 bindings for developing SELinux applications and it is used for backward compatibility. For this reason, **libselinux-python** is no longer available in the default RHEL 8 repositories through the **yum install libselinux-python** command.

To work around this problem, enable both the **libselinux-python** and **python27** modules, and install the **libselinux-python** package and its dependencies with the following commands:

```
# yum module enable libselinux-python
# yum install libselinux-python
```

Alternatively, install **libselinux-python** using its install profile with a single command:

```
# yum module install libselinux-python:2.8/common
```

As a result, you can install **libselinux-python** using the respective module.

Bugzilla:1666328[1]

udica processes UBI 8 containers only when started with --env container=podman

The Red Hat Universal Base Image 8 (UBI 8) containers set the **container** environment variable to the **oci** value instead of the **podman** value. This prevents the **udica** tool from analyzing a container JavaScript Object Notation (JSON) file.

To work around this problem, start a UBI 8 container using a **podman** command with the **--env container=podman** parameter. As a result, **udica** can generate an SELinux policy for a UBI 8 container only when you use the described workaround.

Bugzilla:1763210

Negative effects of the default logging setup on performance

The default logging environment setup might consume 4 GB of memory or even more and adjustments of rate-limit values are complex when **systemd-journald** is running with **rsyslog**.

See the Negative effects of the RHEL default logging setup on performance and their mitigations Knowledgebase article for more information.

Jira:RHELPLAN-10431[1]

SELINUX=disabled in /etc/selinux/config does not work properly

Disabling SELinux using the **SELINUX=disabled** option in the **/etc/selinux/config** results in a process in which the kernel boots with SELinux enabled and switches to disabled mode later in the boot process. This might cause memory leaks.

To work around this problem, disable SELinux by adding the **selinux=0** parameter to the kernel command line as described in the Changing SELinux modes at boot time section of the Using SELinux title if your scenario really requires to completely disable SELinux.

Jira:RHELPLAN-34199[1]

IKE over TCP connections do not work on custom TCP ports

The **tcp-remoteport** Libreswan configuration option does not work properly. Consequently, an IKE over TCP connection cannot be established when a scenario requires specifying a non-default TCP port.

Bugzilla:1989050

scap-security-guide cannot configure termination of idle sessions

Even though the **sshd_set_idle_timeout** rule still exists in the data stream, the former method for idle session timeout of configuring **sshd** is no longer available. Therefore, the rule is marked as **not applicable** and cannot harden anything. Other methods for configuring idle session termination, such as **systemd** (Logind), are also not available. As a consequence, **scap-security-guide** cannot configure the system to reliably disconnect idle sessions after a certain amount of time.

You can work around this problem in one of the following ways, which might fulfill the security requirement:

- Configuring the **accounts_tmout** rule. However, this variable could be overridden by using the **exec** command.

- Configuring the **configure_tmux_lock_after_time** and **configure_bashrc_exec_tmux** rules. This requires installing the **tmux** package.

- Upgrading to RHEL 8.7 or later where the **systemd** feature is already implemented together with the proper SCAP rule.

Jira:RHEL-1804

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-1810

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

Rebuilding the **rpm** database assigns incorrect SELinux labeling

Rebuilding the **rpm** database with the **rpmdb --rebuilddb** command assigns incorrect SELinux labels to the **rpm** database files. As a consequence, some services that use the **rpm** database might not work correctly. To work around this problem after rebuilding the database, relabel the database by using the **restorecon -Rv /var/lib/rpm** command.

Bugzilla:2166153

ANSSI BP28 HP SCAP rules for Audit are incorrectly used on the 64-bit ARM architecture

The ANSSI BP28 High profile in the SCAP Security Guide (SSG) contains the following security content automation protocol (SCAP) rules that configure the Linux Audit subsystem but are invalid on the 64-bit ARM architecture:

- **audit_rules_unsuccessful_file_modification_creat**

- **audit_rules_unsuccessful_file_modification_open**

- **audit_rules_file_deletion_events_rename**

- **audit_rules_file_deletion_events_rmdir**

- **audit_rules_file_deletion_events_unlink**

- **audit_rules_dac_modification_chmod**

- **audit_rules_dac_modification_chown**

- **audit_rules_dac_modification_lchown**

If you configure your RHEL system running on a 64-bit ARM machine by using this profile, the Audit daemon does not start due to the use of invalid system calls.

To work around the problem, either use profile tailoring to remove the previously mentioned rules from the data stream or remove the **-S <syscall>** snippets by editing files in the **/etc/audit/rules.d** directory. The files must not contain the following system calls:

- creat

- open

- rename

- rmdir

- unlink

- chmod

- chown

- lchown

As a result of any of the two described workarounds, the Audit daemon can start even after you use the ANSSI BP28 High profile on a 64-bit ARM system.

Jira:RHEL-1897

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.

Bugzilla:1834716

11.3. SUBSCRIPTION MANAGEMENT

syspurpose addons have no effect on the **subscription-manager attach --auto** output

In Red Hat Enterprise Linux 8, four attributes of the **syspurpose** command-line tool have been added: **role**,**usage**, **service_level_agreement** and **addons**. Currently, only **role**, **usage** and **service_level_agreement** affect the output of running the **subscription-manager attach --auto** command. Users who attempt to set values to the **addons** argument will not observe any effect on the subscriptions that are auto-attached.

Bugzilla:1687900

11.4. SOFTWARE MANAGEMENT

cr_compress_file_with_stat() can cause a memory leak

The **createrepo_c** C library has the API **cr_compress_file_with_stat()** function. This function is declared with **char **dst** as a second parameter. Depending on its other parameters, **cr_compress_file_with_stat()** either uses **dst** as an input parameter, or uses it to return an allocated string. This unpredictable behavior can cause a memory leak, because it does not inform the user when to free **dst** contents.

To work around this problem, a new API **cr_compress_file_with_stat_v2** function has been added, which uses the **dst** parameter only as an input. It is declared as **char *dst**. This prevents memory leak.

Note that the **cr_compress_file_with_stat_v2** function is temporary and will be present only in RHEL 8. Later, **cr_compress_file_with_stat()** will be fixed instead.

Bugzilla:1973588[1]

YUM transactions reported as successful when a scriptlet fails

Since RPM version 4.6, post-install scriptlets are allowed to fail without being fatal to the transaction. This behavior propagates up to YUM as well. This results in scriptlets which might occasionally fail while the overall package transaction reports as successful.

There is no workaround available at the moment.

Note that this is expected behavior that remains consistent between RPM and YUM. Any issues in scriptlets should be addressed at the package level.

Bugzilla:1986657

11.5. SHELLS AND COMMAND-LINE TOOLS

ipmitool is incompatible with certain server platforms

The **ipmitool** utility serves for monitoring, configuring, and managing devices that support the Intelligent Platform Management Interface (IPMI). The current version of **ipmitool** uses Cipher Suite 17 by default instead of the previous Cipher Suite 3. Consequently, **ipmitool** fails to communicate with certain bare metal nodes that announced support for Cipher Suite 17 during negotiation, but do not actually support this cipher suite. As a result, **ipmitool** aborts with the **no matching cipher suite** error message.

For more details, see the related Knowledgebase article.

To solve this problem, update your baseboard management controller (BMC) firmware to use the Cipher Suite 17.

Optionally, if the BMC firmware update is not available, you can work around this problem by forcing **ipmitool** to use a certain cipher suite. When invoking a managing task with **ipmitool**, add the **-C** option to the **ipmitool** command together with the *number* of the cipher suite you want to use. See the following example:

```
# ipmitool -I lanplus -H myserver.example.com -P mypass -C 3 chassis power status
```

Jira:RHEL-6846

ReaR fails to recreate a volume group when you do not use clean disks for restoring

ReaR fails to perform recovery when you want to restore to disks that contain existing data.

To work around this problem, wipe the disks manually before restoring to them if they have been previously used. To wipe the disks in the rescue environment, use one of the following commands before running the **rear recover** command:

- The **dd** command to overwrite the disks.

- The **wipefs** command with the **-a** flag to erase all available metadata.

See the following example of wiping metadata from the **/dev/sda** disk:

```
# wipefs -a /dev/sda[1-9] /dev/sda
```

This command wipes the metadata from the partitions on **/dev/sda** first, and then the partition table itself.

Bugzilla:1925531

coreutils might report misleading EPERM error codes

GNU Core Utilities (**coreutils**) started using the **statx()** system call. If a **seccomp** filter returns an EPERM error code for unknown system calls, **coreutils** might consequently report misleading EPERM error codes because EPERM can not be distinguished from the actual *Operation not permitted* error returned by a working **statx()** syscall.

To work around this problem, update the **seccomp** filter to either permit the **statx()** syscall, or to return an ENOSYS error code for syscalls it does not know.

Bugzilla:2030661

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-12008

11.6. INFRASTRUCTURE SERVICES

Postfix TLS fingerprint algorithm in the FIPS mode needs to be changed to SHA-256

By default in RHEL 8, **postfix** uses MD5 fingerprints with the TLS for backward compatibility. But in the FIPS mode, the MD5 hashing function is not available, which may cause TLS to incorrectly function in the default postfix configuration. To work around this problem, the hashing function needs to be changed to SHA-256 in the postfix configuration file.

For more details, see the related Knowledgebase article Fix postfix TLS in the FIPS mode by switching to SHA-256 instead of MD5.

Bugzilla:1711885

The **brltty** package is not multilib compatible

It is not possible to have both 32-bit and 64-bit versions of the **brltty** package installed. You can either install the 32-bit (**brltty.i686**) or the 64-bit (**brltty.x86_64**) version of the package. The 64-bit version is recommended.

Bugzilla:2008197

11.7. NETWORKING

RoCE interfaces lose their IP settings due to an unexpected change of the network interface name

The RDMA over Converged Ethernet (RoCE) interfaces lose their IP settings due to an unexpected change of the network interface name if both conditions are met:

- User upgrades from a RHEL 8.6 system or earlier.

- The RoCE card is enumerated by UID.

To work around this problem:

1. Create the /**etc**/**systemd**/**network**/**98-rhel87-s390x.link** file with the following content:

   ```
   [Match]
   Architecture=s390x
   KernelCommandLine=!net.naming-scheme=rhel-8.7
   ```

> [Link]
> NamePolicy=kernel database slot path
> AlternativeNamesPolicy=database slot path
> MACAddressPolicy=persistent

2. Reboot the system for the changes to take effect.

3. Upgrade to RHEL 8.7 or newer.

Note that RoCE interfaces that are enumerated by function ID (FID) and are non-unique, will still use unpredictable interface names unless you set the **net.naming-scheme=rhel-8.7** kernel parameter. In this case, the RoCE interfaces will switch to predictable names with the **ens** prefix.

Jira:RHEL-11398[1]

Systems with the **IPv6_rpfilter** option enabled experience low network throughput

Systems with the **IPv6_rpfilter** option enabled in the **firewalld.conf** file currently experience suboptimal performance and low network throughput in high traffic scenarios, such as 100 Gbps links. To work around the problem, disable the **IPv6_rpfilter** option. To do so, add the following line in the **/etc/firewalld/firewalld.conf** file.

> IPv6_rpfilter=no

As a result, the system performs better, but also has reduced security.

Bugzilla:1871860[1]

11.8. KERNEL

The kernel ACPI driver reports it has no access to a PCIe ECAM memory region

The Advanced Configuration and Power Interface (ACPI) table provided by firmware does not define a memory region on the PCI bus in the Current Resource Settings (_CRS) method for the PCI bus device. Consequently, the following warning message occurs during the system boot:

> [2.817152] acpi PNP0A08:00: [Firmware Bug]: ECAM area [mem 0x30000000-0x31ffffff] not reserved in ACPI namespace
> [2.827911] acpi PNP0A08:00: ECAM at [mem 0x30000000-0x31ffffff] for [bus 00-1f]

However, the kernel is still able to access the **0x30000000-0x31ffffff** memory region, and can assign that memory region to the PCI Enhanced Configuration Access Mechanism (ECAM) properly. You can verify that PCI ECAM works correctly by accessing the PCIe configuration space over the 256 byte offset with the following output:

> 03:00.0 Non-Volatile memory controller: Sandisk Corp WD Black 2018/PC SN720 NVMe SSD (prog-if 02 [NVM Express])
> ...
>   Capabilities: [900 v1] L1 PM Substates
>   L1SubCap: PCI-PM_L1.2- PCI-PM_L1.1- ASPM_L1.2+ ASPM_L1.1- L1_PM_Substates+
>   PortCommonModeRestoreTime=255us PortTPowerOnTime=10us
>   L1SubCtl1: PCI-PM_L1.2- PCI-PM_L1.1- ASPM_L1.2- ASPM_L1.1-
>   T_CommonMode=0us LTR1.2_Threshold=0ns
>   L1SubCtl2: T_PwrOn=10us

As a result, you can ignore the warning message.

For more information about the problem, see the "Firmware Bug: ECAM area **mem 0x30000000-0x31ffffff** not reserved in ACPI namespace" appears during system boot solution.

Bugzilla:1868526[1]

The **tuned-adm profile powersave** command causes the system to become unresponsive

Executing the **tuned-adm profile powersave** command leads to an unresponsive state of the Penguin Valkyrie 2000 2-socket systems with the older Thunderx (CN88xx) processors. Consequently, reboot the system to resume working. To work around this problem, avoid using the **powersave** profile if your system matches the mentioned specifications.

Bugzilla:1609288[1]

The HP NMI watchdog does not always generate a crash dump

In certain cases, the **hpwdt** driver for the HP NMI watchdog is not able to claim a non-maskable interrupt (NMI) generated by the HPE watchdog timer because the NMI was instead consumed by the **perfmon** driver.

The missing NMI is initiated by one of two conditions:

1. The **Generate NMI** button on the Integrated Lights-Out (iLO) server management software. This button is triggered by a user.

2. The **hpwdt** watchdog. The expiration by default sends an NMI to the server.

Both sequences typically occur when the system is unresponsive. Under normal circumstances, the NMI handler for both these situations calls the **kernel panic()** function and if configured, the **kdump** service generates a **vmcore** file.

Because of the missing NMI, however, **kernel panic()** is not called and **vmcore** is not collected.

In the first case (1.), if the system was unresponsive, it remains so. To work around this scenario, use the virtual **Power** button to reset or power cycle the server.

In the second case (2.), the missing NMI is followed 9 seconds later by a reset from the Automated System Recovery (ASR).

The HPE Gen9 Server line experiences this problem in single-digit percentages. The Gen10 at an even smaller frequency.

Bugzilla:1602962[1]

Reloading an identical crash extension may cause segmentation faults

When you load a copy of an already loaded crash extension file, it might trigger a segmentation fault. Currently, the crash utility detects if an original file has been loaded. Consequently, due to two identical files co-existing in the crash utility, a namespace collision occurs, which triggers the crash utility to cause a segmentation fault.

You can work around the problem by loading the crash extension file only once. As a result, segmentation faults no longer occur in the described scenario.

Bugzilla:1906482

Connections fail when attaching a virtual function to virtual machine

Pensando network cards that use the **ionic** device driver silently accept VLAN tag configuration requests and attempt configuring network connections while attaching network virtual functions (**VF**) to a virtual machine (**VM**). Such network connections fail as this feature is not yet supported by the card's firmware.

Bugzilla:1930576[1]

The OPEN MPI library may trigger run-time failures with default PML

In OPEN Message Passing Interface (OPEN MPI) implementation 4.0.x series, Unified Communication X (UCX) is the default point-to-point communicator (PML). The later versions of OPEN MPI 4.0.x series deprecated **openib** Byte Transfer Layer (BTL).

However, OPEN MPI, when run over a **homogeneous** cluster (same hardware and software configuration), UCX still uses **openib** BTL for MPI one-sided operations. As a consequence, this may trigger execution errors. To work around this problem:

- Run the **mpirun** command using following parameters:

```
-mca btl openib -mca pml ucx -x UCX_NET_DEVICES=mlx5_ib0
```

where,

- The **-mca btl openib** parameter disables **openib** BTL

- The **-mca pml ucx** parameter configures OPEN MPI to use **ucx** PML.

- The **x UCX_NET_DEVICES=** parameter restricts UCX to use the specified devices

The OPEN MPI, when run over a **heterogeneous** cluster (different hardware and software configuration), it uses UCX as the default PML. As a consequence, this may cause the OPEN MPI jobs to run with erratic performance, unresponsive behavior, or crash failures. To work around this problem, set the UCX priority as:

- Run the **mpirun** command using following parameters:

```
-mca pml_ucx_priority 5
```

As a result, the OPEN MPI library is able to choose an alternative available transport layer over UCX.

Bugzilla:1866402[1]

vmcore capture fails after memory hot-plug or unplug operation

After performing the memory hot-plug or hot-unplug operation, the event comes after updating the device tree which contains memory layout information. Thereby the **makedumpfile** utility tries to access a non-existent physical address. The problem appears if all of the following conditions meet:

- A little-endian variant of IBM Power System runs RHEL 8.

- The **kdump** or **fadump** service is enabled on the system.

Consequently, the capture kernel fails to save **vmcore** if a kernel crash is triggered after the memory hot-plug or hot-unplug operation.

To work around this problem, restart the **kdump** service after hot-plug or hot-unplug:

```
# systemctl restart kdump.service
```

As a result, **vmcore** is successfully saved in the described scenario.

Bugzilla:1793389[1]

Using **irqpoll** causes **vmcore** generation failure

Due to an existing problem with the **nvme** driver on the 64-bit ARM architecture that run on the Amazon Web Services Graviton 1 processor, causes **vmcore** generation to fail when you provide the **irqpoll** kernel command line parameter to the first kernel. Consequently, no **vmcore** file is dumped in the **/var/crash/** directory upon a kernel crash. To work around this problem:

1. Append **irqpoll** to **KDUMP_COMMANDLINE_REMOVE** variable in the **/etc/sysconfig/kdump** file.

   ```
   # KDUMP_COMMANDLINE_REMOVE="hugepages hugepagesz slub_debug quiet
   log_buf_len swiotlb"
   ```

2. Remove **irqpoll** from **KDUMP_COMMANDLINE_APPEND** variable in the **/etc/sysconfig/kdump** file.

   ```
   # KDUMP_COMMANDLINE_APPEND="irqpoll nr_cpus=1 reset_devices
   cgroup_disable=memory udev.children-max=2 panic=10 swiotlb=noforce novmcoredd"
   ```

3. Restart the **kdump** service:

   ```
   # systemctl restart kdump
   ```

As a result, the first kernel boots correctly and the **vmcore** file is expected to be captured upon the kernel crash.

Note that the Amazon Web Services Graviton 2 and Amazon Web Services Graviton 3 processors do not require you to manually remove the **irqpoll** parameter in the **/etc/sysconfig/kdump** file.

The **kdump** service can use a significant amount of crash kernel memory to dump the **vmcore** file. Ensure that the capture kernel has sufficient memory available for the **kdump** service.

For related information on this Known Issue, see The irqpoll kernel command line parameter might cause vmcore generation failure article.

Bugzilla:1654962[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]

Debug kernel fails to boot in crash capture environment on RHEL 8

Due to the memory-intensive nature of the debug kernel, a problem occurs when the debug kernel is in use and a kernel panic is triggered. As a consequence, the debug kernel is not able to boot as the capture kernel and a stack trace is generated instead. To work around this problem, increase the crash kernel memory as required. As a result, the debug kernel boots successfully in the crash capture environment.

Bugzilla:1659609[1]

Allocating crash kernel memory fails at boot time

On some Ampere Altra systems, allocating the crash kernel memory during boot fails when the 32-bit region is disabled in BIOS settings. Consequently, the **kdump** service fails to start. This is caused by memory fragmentation in the region below 4 GB with no fragment being large enough to contain the crash kernel memory.

To work around this problem, enable the 32-bit memory region in BIOS as follows:

1. Open the BIOS settings on your system.

2. Open the **Chipset** menu.

3. Under **Memory Configuration**, enable the **Slave 32-bit** option.

As a result, crash kernel memory allocation within the 32-bit region succeeds and the **kdump** service works as expected.

Bugzilla:1940674[1]

The QAT manager leaves no spare device for LKCF

The Intel® QuickAssist Technology (QAT) manager (**qatmgr**) is a user space process, which by default uses all QAT devices in the system. As a consequence, there are no QAT devices left for the Linux Kernel Cryptographic Framework (LKCF). There is no need to work around this situation, as this behavior is expected and a majority of users will use acceleration from the user space.

Bugzilla:1920086[1]

The Solarflare fails to create maximum number of virtual functions (VFs)

The Solarflare NICs fail to create a maximum number of VFs due to insufficient resources. You can check

the maximum number of VFs that a PCIe device can create in the
/sys/bus/pci/devices/PCI_ID/sriov_totalvfs file. To workaround this problem, you can either adjust the
number of VFs or the VF MSI interrupt value to a lower value, either from **Solarflare Boot Manager** on
startup, or using Solarflare **sfboot** utility. The default VF MSI interrupt value is 8.

- To adjust the VF MSI interrupt value using **sfboot**:

```
# sfboot vf-msix-limit=2
```
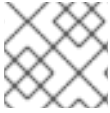
> **NOTE**
>
> Adjusting VF MSI interrupt value affects the VF performance.

For more information about parameters to be adjusted accordingly, see the **Solarflare Server Adapter
user guide**.

Bugzilla:1971506[1]

Using page_poison=1 can cause a kernel crash

When using **page_poison=1** as the kernel parameter on firmware with faulty EFI implementation, the
operating system can cause the kernel to crash. By default, this option is disabled and it is not
recommended to enable it, especially in production systems.

Bugzilla:2050411[1]

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 8.7 and
later, the hardware gets into an incorrect internal state. reports its state incorrectly. Consequently, Intel
Wifi 6 cards may not work and display the error message:

```
kernel: iwlwifi 0000:09:00.0: Failed to start RT ucode: -110
kernel: iwlwifi 0000:09:00.0: WRT: Collecting data: ini trigger 13 fired (delay=0ms)
kernel: iwlwifi 0000:09:00.0: Failed to run INIT ucode: -110
```

An unconfirmed work around is to power off the system and back on again. Do not reboot.

Bugzilla:2106341[1]

Secure boot on IBM Power Systems does not support migration

Currently, on IBM Power Systems, logical partition (LPAR) does not boot after successful physical
volume (PV) migration. As a result, any type of automated migration with secure boot enabled on a
partition fails.

Bugzilla:2126777[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]

kdump in Ampere Altra servers enters the OOM state

The firmware in Ampere Altra and Altra Max servers currently causes the kernel to allocate too many event, interrupt and command queues, which consumes too much memory. As a consequence, the **kdump** kernel enters the Out of memory (OOM) state.

To work around this problem, reserve extra memory for **kdump** by increasing the value of the **crashkernel=** kernel option to *640M*.

Bugzilla:2111855[1]

11.9. FILE SYSTEMS AND STORAGE

LVM **mirror** devices that store a LUKS volume sometimes become unresponsive

Mirrored LVM devices with a segment type of **mirror** that store a LUKS volume might become unresponsive under certain conditions. The unresponsive devices reject all I/O operations.

To work around the issue, Red Hat recommends that you use LVM RAID 1 devices with a segment type of **raid1** instead of **mirror** if you need to stack LUKS volumes on top of resilient software-defined storage.

The **raid1** segment type is the default RAID configuration type and replaces **mirror** as the recommended solution.

To convert **mirror** devices to **raid1**, see Converting a mirrored LVM device to a RAID1 device .

Bugzilla:1730502[1]

The /**boot** file system cannot be placed on LVM

You cannot place the /**boot** file system on an LVM logical volume. This limitation exists for the following reasons:

- On EFI systems, the *EFI System Partition* conventionally serves as the /**boot** file system. The uEFI standard requires a specific GPT partition type and a specific file system type for this partition.

- RHEL 8 uses the *Boot Loader Specification* (BLS) for system boot entries. This specification requires that the /**boot** file system is readable by the platform firmware. On EFI systems, the platform firmware can read only the /**boot** configuration defined by the uEFI standard.

- The support for LVM logical volumes in the GRUB 2 boot loader is incomplete. Red Hat does not plan to improve the support because the number of use cases for the feature is decreasing due to standards such as uEFI and BLS.

Red Hat does not plan to support /**boot** on LVM. Instead, Red Hat provides tools for managing system snapshots and rollback that do not need the /**boot** file system to be placed on an LVM logical volume.

Bugzilla:1496229[1]

LVM no longer allows creating volume groups with mixed block sizes

LVM utilities such as **vgcreate** or **vgextend** no longer allow you to create volume groups (VGs) where the physical volumes (PVs) have different logical block sizes. LVM has adopted this change because file systems fail to mount if you extend the underlying logical volume (LV) with a PV of a different block size.

To re-enable creating VGs with mixed block sizes, set the **allow_mixed_block_sizes=1** option in the **lvm.conf** file.

Bugzilla:1768536

Limitations of LVM **writecache**

The **writecache** LVM caching method has the following limitations, which are not present in the **cache** method:

- You cannot name a **writecache** logical volume when using **pvmove** commands.

- You cannot use logical volumes with **writecache** in combination with thin pools or VDO.

The following limitation also applies to the **cache** method:

- You cannot resize a logical volume while **cache** or **writecache** is attached to it.

Jira:RHELPLAN-27987[1], Bugzilla:1808012, Bugzilla:1798631

Device-mapper multipath is not supported when using NVMe/TCP driver.

The use of device-mapper multipath on top of NVMe/TCP devices can cause reduced performance and error handling. To avoid this problem, use native NVMe multipath instead of DM multipath tools. For RHEL 8, you can add the option **nvme_core.multipath=Y** to the kernel command line.

Bugzilla:2022359[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]

XFS quota warnings are triggered too often

Using the quota timer results in quota warnings triggering too often, which causes soft quotas to be enforced faster than they should. To work around this problem, do not use soft quotas, which will prevent triggering warnings. As a result, the amount of warning messages will not enforce soft quota limit anymore, respecting the configured timeout.

Bugzilla:2059262[1]

11.10. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Creating virtual Python 3.11 environments fails when using the virtualenv utility

The **virtualenv** utility in RHEL 8, provided by the **python3-virtualenv** package, is not compatible with Python 3.11. An attempt to create a virtual environment by using **virtualenv** will fail with the following error message:

```
$ virtualenv -p python3.11 venv3.11
Running virtualenv with interpreter /usr/bin/python3.11
ERROR: Virtual environments created by virtualenv < 20 are not compatible with Python 3.11.
ERROR: Use `python3.11 -m venv` instead.
```

To create Python 3.11 virtual environments, use the **python3.11 -m venv** command instead, which uses the **venv** module from the standard library.

Bugzilla:2165702

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:2157673

PAM plug-in version 1.0 does not work in MariaDB

MariaDB 10.3 provides the Pluggable Authentication Modules (PAM) plug-in version 1.0. **MariaDB 10.5** provides the plug-in versions 1.0 and 2.0, version 2.0 is the default.

The **MariaDB** PAM plug-in version 1.0 does not work in RHEL 8. To work around this problem, use the PAM plug-in version 2.0 provided by the **mariadb:10.5** module stream.

Bugzilla:1942330

Symbol conflicts between OpenLDAP libraries might cause crashes in httpd

When both the **libldap** and **libldap_r** libraries provided by OpenLDAP are loaded and used within a single process, symbol conflicts between these libraries might occur. Consequently, Apache **httpd** child processes using the PHP **ldap** extension might terminate unexpectedly if the **mod_security** or **mod_auth_openidc** modules are also loaded by the **httpd** configuration.

Since the RHEL 8.3 update to the Apache Portable Runtime (APR) library, you can work around the problem by setting the **APR_DEEPBIND** environment variable, which enables the use of the **RTLD_DEEPBIND** dynamic linker option when loading **httpd** modules. When the **APR_DEEPBIND** environment variable is enabled, crashes no longer occur in **httpd** configurations that load conflicting libraries.

Bugzilla:1819607[1]

getpwnam() might fail when called by a 32-bit application

When a user of NIS uses a 32-bit application that calls the **getpwnam()** function, the call fails if the **nss_nis.i686** package is missing. To work around this problem, manually install the missing package by using the **yum install nss_nis.i686** command.

Bugzilla:1803161

11.11. IDENTITY MANAGEMENT

Actions required when running Samba as a print server and updating from RHEL 8.4 and earlier

With this update, the **samba** package no longer creates the /**var**/**spool**/**samba**/ directory. If you use Samba as a print server and use /**var**/**spool**/**samba**/ in the **[printers]** share to spool print jobs, SELinux prevents Samba users from creating files in this directory. Consequently, print jobs fail and the **auditd** service logs a **denied** message in /**var**/**log**/**audit**/**audit.log**. To avoid this problem after updating your system from 8.4 and earlier:

1. Search the **[printers]** share in the /**etc**/**samba**/**smb.conf** file.

2. If the share definition contains **path = /var/spool/samba/**, update the setting and set the **path** parameter to /**var**/**tmp**/.

3. Restart the **smbd** service:

   ```
   # systemctl restart smbd
   ```

If you newly installed Samba on RHEL 8.5 or later, no action is required. The default /**etc**/**samba**/**smb.conf** file provided by the **samba-common** package in this case already uses the /**var**/**tmp**/ directory to spool print jobs.

Bugzilla:2009213[1]

Using the `cert-fix` utility with the `--agent-uid pkidbuser` option breaks Certificate System

Using the **cert-fix** utility with the **--agent-uid pkidbuser** option corrupts the LDAP configuration of Certificate System. As a consequence, Certificate System might become unstable and manual steps are required to recover the system.

Bugzilla:1729215

FIPS mode does not support using a shared secret to establish a cross-forest trust

Establishing a cross-forest trust using a shared secret fails in FIPS mode because NTLMSSP authentication is not FIPS-compliant. To work around this problem, authenticate with an Active Directory (AD) administrative account when establishing a trust between an IdM domain with FIPS mode enabled and an AD domain.

Jira:RHEL-4847

Downgrading `authselect` after the rebase to version 1.2.2 breaks system authentication

The **authselect** package has been rebased to the latest upstream version **1.2.2**. Downgrading **authselect** is not supported and breaks system authentication for all users, including **root**.

If you downgraded the **authselect** package to **1.2.1** or earlier, perform the following steps to work around this problem:

1. At the GRUB boot screen, select **Red Hat Enterprise Linux** with the version of the kernel that you want to boot and press **e** to edit the entry.

2. Type **single** as a separate word at the end of the line that starts with **linux** and press **Ctrl+X** to start the boot process.

3. Upon booting in single-user mode, enter the root password.

4. Restore authselect configuration using the following command:

   ```
   # authselect select sssd --force
   ```

Bugzilla:1892761

IdM to AD cross-realm TGS requests fail

The Privilege Attribute Certificate (PAC) information in IdM Kerberos tickets is now signed with AES SHA-2 HMAC encryption, which is not supported by Active Directory (AD).

Consequently, IdM to AD cross-realm TGS requests, that is, two-way trust setups, are failing with the following error:

```
Generic error (see e-text) while getting credentials for <service principal>
```

Jira:RHEL-4910

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]

pki-core-debuginfo update from RHEL 8.6 to RHEL 8.7 or later fails

Updating the **pki-core-debuginfo** package from RHEL 8.6 to RHEL 8.7 or later fails. To work around this problem, run the following commands:

1. **yum remove pki-core-debuginfo**

2. **yum update -y**

3. **yum install pki-core-debuginfo**

4. **yum install idm-pki-symkey-debuginfo idm-pki-tools-debuginfo**

Jira:RHEL-13125[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]

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-4898

IdM Vault encryption and decryption fails in FIPS mode

The OpenSSL RSA-PKCS1v15 padding encryption is blocked if FIPS mode is enabled. Consequently, Identity Management (IdM) Vaults fail to work correctly as IdM is currently using the PKCS1v15 padding for wrapping the session key with the transport certificate.

Jira:RHEL-12153[1]

Incorrect warning when setting expiration dates for a Kerberos principal

If you set a password expiration date for a Kerberos principal, the current timestamp is compared to the expiration timestamp using a 32-bit signed integer variable. If the expiration date is more than 68 years in the future, it causes an integer variable overflow resulting in the following warning message being displayed:

> Warning: Your password will expire in less than one hour on [expiration date]

You can ignore this message, the password will expire correctly at the configured date and time.

Bugzilla:2125318

11.12. DESKTOP

Disabling flatpak repositories from Software Repositories is not possible

Currently, it is not possible to disable or remove **flatpak** repositories in the Software Repositories tool in the GNOME Software utility.

Bugzilla:1668760

Generation 2 RHEL 8 virtual machines sometimes fail to boot on Hyper-V Server 2016 hosts

When using RHEL 8 as the guest operating system on a virtual machine (VM) running on a Microsoft Hyper-V Server 2016 host, the VM in some cases fails to boot and returns to the GRUB boot menu. In addition, the following error is logged in the Hyper-V event log:

> The guest operating system reported that it failed with the following error code: 0x1E

This error occurs due to a UEFI firmware bug on the Hyper-V host. To work around this problem, use Hyper-V Server 2019 or later as the host.

Bugzilla:1583445[1]

Drag-and-drop does not work between desktop and applications

Due to a bug in the **gnome-shell-extensions** package, the drag-and-drop functionality does not currently work between desktop and applications. Support for this feature will be added back in a future release.

Bugzilla:1717947

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 terminates 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-4158

11.13. GRAPHICS INFRASTRUCTURES

The **radeon** driver fails to reset hardware correctly

The **radeon** kernel driver currently does not reset hardware in the **kexec** context correctly. Instead, **radeon** falls over, which causes the rest of the **kdump** service to fail.

To work around this problem, disable **radeon** in **kdump** by adding the following line to the **/etc/kdump.conf** file:

```
dracut_args --omit-drivers "radeon"
force_rebuild 1
```

Restart the system and **kdump**. After starting **kdump**, the **force_rebuild 1** line might be removed from the configuration file.

Note that in this scenario, no graphics is available during the dump process, but **kdump** works correctly.

Bugzilla:1694705[1]

Multiple HDR displays on a single MST topology may not power on

On systems using NVIDIA Turing GPUs with the **nouveau** driver, using a **DisplayPort** hub (such as a laptop dock) with multiple monitors which support HDR plugged into it may result in failure to turn on. This is due to the system erroneously thinking there is not enough bandwidth on the hub to support all

of the displays.

Bugzilla:1812577[1]

GUI in ESXi might crash due to low video memory

The graphical user interface (GUI) on RHEL virtual machines (VMs) in the VMware ESXi 7.0.1 hypervisor with vCenter Server 7.0.1 requires a certain amount of video memory. If you connect multiple consoles or high-resolution monitors to the VM, the GUI requires at least 16 MB of video memory. If you start the GUI with less video memory, the GUI might terminate unexpectedly.

To work around the problem, configure the hypervisor to assign at least 16 MB of video memory to the VM. As a result, the GUI on the VM no longer crashes.

If you encounter this issue, Red Hat recommends that you report it to VMware.

See also the following VMware article: VMs with high resolution VM console may experience a crash on ESXi 7.0.1 (83194).

Bugzilla:1910358[1]

VNC Viewer displays wrong colors with the 16-bit color depth on IBM Z

The VNC Viewer application displays wrong colors when you connect to a VNC session on an IBM Z server with the 16-bit color depth.

To work around the problem, set the 24-bit color depth on the VNC server. With the **Xvnc** server, replace the **-depth 16** option with **-depth 24** in the **Xvnc** configuration.

As a result, VNC clients display the correct colors but use more network bandwidth with the server.

Bugzilla:1886147

Unable to run graphical applications using **sudo** command

When trying to run graphical applications as a user with elevated privileges, the application fails to open with an error message. The failure happens because **Xwayland** is restricted by the **Xauthority** file to use regular user credentials for authentication.

To work around this problem, use the **sudo -E** command to run graphical applications as a **root** user.

Bugzilla:1673073

Hardware acceleration is not supported on ARM

Built-in graphics drivers do not support hardware acceleration or the Vulkan API on the 64-bit ARM architecture.

To enable hardware acceleration or Vulkan on ARM, install the proprietary Nvidia driver.

Jira:RHELPLAN-57914[1]

11.14. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Unable to manage **localhost** by using the **localhost** hostname in the playbook or inventory

With the inclusion of the **ansible-core 2.13** package in RHEL, if you are running Ansible on the same host you manage your nodes, you cannot do it by using the **localhost** hostname in your playbook or

inventory. This happens because **ansible-core 2.13** uses the **python38** module, and many of the libraries are missing, for example, **blivet** for the **storage** role, **gobject** for the **network** role. To workaround this problem, if you are already using the **localhost** hostname in your playbook or inventory, you can add a connection, by using **ansible_connection=local**, or by creating an inventory file that lists **localhost** with the **ansible_connection=local** option. With that, you are able to manage resources on **localhost**. For more details, see the article RHEL System Roles playbooks fail when run on localhost .

Bugzilla:2041997

The **rhc** system role fails on already registered systems when **rhc_auth** contains activation keys

Executing playbook files on already registered systems fails if activation keys are specified for the **rhc_auth** parameter. To workaround this issue, do not specify activation keys when executing the playbook file on the already registered system.

Bugzilla:2186908

11.15. VIRTUALIZATION

Using a large number of queues might cause Windows virtual machines to fail

Windows virtual machines (VMs) might fail when the virtual Trusted Platform Module (vTPM) device is enabled and the *multi-queue virtio-net* feature is configured to use more than 250 queues.

This problem is caused by a limitation in the vTPM device. The vTPM device has a hardcoded limit on the maximum number of opened file descriptors. Since multiple file descriptors are opened for every new queue, the internal vTPM limit can be exceeded, causing the VM to fail.

To work around this problem, choose one of the following two options:

- Keep the vTPM device enabled, but use less than 250 queues.

- Disable the vTPM device to use more than 250 queues.

Jira:RHEL-13336[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:2077770[1]

SMT CPU topology is not detected by VMs when using host passthrough mode on AMD EPYC

When a virtual machine (VM) boots with the CPU host passthrough mode on an AMD EPYC host, the **TOPOEXT** CPU feature flag is not present. Consequently, the VM is not able to detect a virtual CPU topology with multiple threads per core. To work around this problem, boot the VM with the EPYC CPU model instead of host passthrough.

Bugzilla:1740002

Attaching LUN devices to virtual machines using virtio-blk does not work

The q35 machine type does not support transitional virtio 1.0 devices, and RHEL 8 therefore lacks support for features that were deprecated in virtio 1.0. In particular, it is not possible on a RHEL 8 host to send SCSI commands from virtio-blk devices. As a consequence, attaching a physical disk as a LUN device to a virtual machine fails when using the virtio-blk controller.

Note that physical disks can still be passed through to the guest operating system, but they should be configured with the **device='disk'** option rather than **device='lun'**.

Bugzilla:1777138[1]

Virtual machines sometimes fail to start when using many virtio-blk disks

Adding a large number of virtio-blk devices to a virtual machine (VM) may exhaust the number of interrupt vectors available in the platform. If this occurs, the VM's guest OS fails to boot, and displays a **dracut-initqueue[392]: Warning: Could not boot** error.

Bugzilla:1719687

Virtual machines with **iommu_platform=on** fail to start on IBM POWER

RHEL 8 currently does not support the **iommu_platform=on** parameter for virtual machines (VMs) on IBM POWER system. As a consequence, starting a VM with this parameter on IBM POWER hardware results in the VM becoming unresponsive during the boot process.

Bugzilla:1910848

IBM POWER hosts now work correctly when using the **ibmvfc** driver

When running RHEL 8 on a PowerVM logical partition (LPAR), a variety of errors could previously occur due to problems with the **ibmvfc** driver. As a consequence, a kernel panic triggered on the host under certain circumstances, such as:

- Using the Live Partition Mobility (LPM) feature

- Resetting a host adapter

- Using SCSI error handling (SCSI EH) functions

With this update, the handling of **ibmvfc** has been fixed, and the described kernel panics no longer occur.

Bugzilla:1961722[1]

Using **perf kvm record** on IBM POWER Systems can cause the VM to crash

When using a RHEL 8 host on the little-endian variant of IBM POWER hardware, using the **perf kvm record** command to collect trace event samples for a KVM virtual machine (VM) in some cases results in the VM becoming unresponsive. This situation occurs when:

- The **perf** utility is used by an unprivileged user, and the **-p** option is used to identify the VM - for example **perf kvm record -e trace_cycles -p 12345**.

- The VM was started using the **virsh** shell.

To work around this problem, use the **perf kvm** utility with the **-i** option to monitor VMs that were created using the **virsh** shell. For example:

> # perf kvm record -e trace_imc/trace_cycles/ -p <guest pid> -i

Note that when using the **-i** option, child tasks do not inherit counters, and threads will therefore not be monitored.

Bugzilla:1924016[1]

Windows Server 2016 virtual machines with Hyper-V enabled fail to boot when using certain CPU models

Currently, it is not possible to boot a virtual machine (VM) that uses Windows Server 2016 as the guest operating system, has the Hyper-V role enabled, and uses one of the following CPU models:

- EPYC-IBPB

- EPYC

To work around this problem, use the **EPYC-v3** CPU model, or manually enable the **xsaves** CPU flag for the VM.

Bugzilla:1942888[1]

Migrating a POWER9 guest from a RHEL 7-ALT host to RHEL 8 fails

Currently, migrating a POWER9 virtual machine from a RHEL 7-ALT host system to RHEL 8 becomes unresponsive with a **Migration status: active** status.

To work around this problem, disable Transparent Huge Pages (THP) on the RHEL 7-ALT host, which enables the migration to complete successfully.

Bugzilla:1741436[1]

Using virt-customize sometimes causes guestfs-firstboot to fail

After modifying a virtual machine (VM) disk image using the **virt-customize** utility, the **guestfs-firstboot** service in some cases fails due to incorrect SELinux permissions. This causes a variety of problems during VM startup, such as failing user creation or system registration.

To avoid this problem, use the **virt-customize** command with the **--selinux-relabel** option.

Bugzilla:1554735

Deleting a forward interface from a macvtap virtual network resets all connection counts of this network

Currently, deleting a forward interface from a **macvtap** virtual network with multiple forward interfaces also resets the connection status of the other forward interfaces of the network. As a consequence, the connection information in the live network XML is incorrect. Note, however, that this does not affect the functionality of the virtual network. To work around the issue, restart the **libvirtd** service on your host.

Bugzilla:1332758

Virtual machines with SLOF fail to boot in netcat interfaces

When using a netcat (**nc**) interface to access the console of a virtual machine (VM) that is currently waiting at the Slimline Open Firmware (SLOF) prompt, the user input is ignored and VM stays unresponsive. To work around this problem, use the **nc -C** option when connecting to the VM, or use a telnet interface instead.

Bugzilla:1974622[1]

Attaching mediated devices to virtual machines in virt-manager in some cases fails

The **virt-manager** application is currently able to detect mediated devices, but cannot recognize whether the device is active. As a consequence, attempting to attach an inactive mediated device to a running virtual machine (VM) using **virt-manager** fails. Similarly, attempting to create a new VM that uses an inactive mediated device fails with a **device not found** error.

To work around this issue, use the **virsh nodedev-start** or **mdevctl start** commands to activate the mediated device before using it in **virt-manager**.

Bugzilla:2026985

RHEL 9 virtual machines fail to boot in POWER8 compatibility mode

Currently, booting a virtual machine (VM) that runs RHEL 9 as its guest operating system fails if the VM also uses CPU configuration similar to the following:

```
<cpu mode="host-model">
  <model>power8</model>
</cpu>
```

To work around this problem, do not use POWER8 compatibility mode in RHEL 9 VMs.

In addition, note that running RHEL 9 VMs is not possible on POWER8 hosts.

Bugzilla:2035158

SUID and SGID are not cleared automatically on virtiofs

When you run the **virtiofsd** service with the **killpriv_v2** feature, your system may not automatically clear the SUID and SGID permissions after performing some file-system operations. Consequently, not clearing the permissions might cause a potential security threat. To work around this issue, disable the **killpriv_v2** feature by entering the following command:

```
# virtiofsd -o no_killpriv_v2
```

Bugzilla:1966475[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.

Bugzilla:1792683

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:2177957

nodedev-dumpxml does not list attributes correctly for certain mediated devices

Currently, the **nodedev-dumpxml** does not list attributes correctly for mediated devices that were created using the **nodedev-create** command. To work around this problem, use the **nodedev-define** and **nodedev-start** commands instead.

Bugzilla:2143160

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-2451[1]

11.16. RHEL IN CLOUD ENVIRONMENTS

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

kdump sometimes does not start on Azure and Hyper-V

On RHEL 8 guest operating systems hosted on the Microsoft Azure or Hyper-V hypervisors, starting the **kdump** kernel in some cases fails when post-exec notifiers are enabled.

To work around this problem, disable crash kexec post notifiers:

```
# echo N > /sys/module/kernel/parameters/crash_kexec_post_notifiers
```

Bugzilla:1865745[1]

The SCSI host address sometimes changes when booting a Hyper-V VM with multiple guest disks

Currently, when booting a RHEL 8 virtual machine (VM) on the Hyper-V hypervisor, the host portion of the *Host, Bus, Target, Lun* (HBTL) SCSI address in some cases changes. As a consequence, automated tasks set up with the HBTL SCSI identification or device node in the VM do not work consistently. This occurs if the VM has more than one disk or if the disks have different sizes.

To work around the problem, modify your kickstart files, using one of the following methods:

Method 1: Use persistent identifiers for SCSI devices.

You can use for example the following powershell script to determine the specific device identifiers:

```
# Output what the /dev/disk/by-id/<value> for the specified hyper-v virtual disk.
# Takes a single parameter which is the virtual disk file.
# Note: kickstart syntax works with and without the /dev/ prefix.
param (
    [Parameter(Mandatory=$true)][string]$virtualdisk
)

$what = Get-VHD -Path $virtualdisk
$part = $what.DiskIdentifier.ToLower().split('-')

$p = $part[0]
$s0 = $p[6] + $p[7] + $p[4] + $p[5] + $p[2] + $p[3] + $p[0] + $p[1]

$p = $part[1]
$s1 =  $p[2] + $p[3] + $p[0] + $p[1]

[string]::format("/dev/disk/by-id/wwn-0x60022480{0}{1}{2}", $s0, $s1, $part[4])
```

You can use this script on the hyper-v host, for example as follows:

```
PS C:\Users\Public\Documents\Hyper-V\Virtual hard disks> .\by-id.ps1 .\Testing_8\disk_3_8.vhdx
/dev/disk/by-id/wwn-0x60022480e00bc367d7fd902e8bf0d3b4
PS C:\Users\Public\Documents\Hyper-V\Virtual hard disks> .\by-id.ps1 .\Testing_8\disk_3_9.vhdx
/dev/disk/by-id/wwn-0x600224807270e09717645b1890f8a9a2
```

Afterwards, the disk values can be used in the kickstart file, for example as follows:

```
part / --fstype=xfs --grow --asprimary --size=8192 --ondisk=/dev/disk/by-id/wwn-
0x600224807270e09717645b1890f8a9a2
part /home --fstype="xfs" --grow --ondisk=/dev/disk/by-id/wwn-
0x60022480e00bc367d7fd902e8bf0d3b4
```

As these values are specific for each virtual disk, the configuration needs to be done for each VM instance. It may, therefore, be useful to use the **%include** syntax to place the disk information into a separate file.

Method 2: Set up device selection by size.

A kickstart file that configures disk selection based on size must include lines similar to the following:

```
...

# Disk partitioning information is supplied in a file to kick start
%include /tmp/disks
```

```
...

# Partition information is created during install using the %pre section
%pre --interpreter /bin/bash --log /tmp/ks_pre.log

 # Dump whole SCSI/IDE disks out sorted from smallest to largest ouputting
 # just the name
 disks=(`lsblk -n -o NAME -l -b -x SIZE -d -I 8,3`) || exit 1

 # We are assuming we have 3 disks which will be used
 # and we will create some variables to represent
 d0=${disks[0]}
 d1=${disks[1]}
 d2=${disks[2]}

 echo "part /home --fstype="xfs" --ondisk=$d2 --grow" >> /tmp/disks
 echo "part swap --fstype="swap" --ondisk=$d0 --size=4096" >> /tmp/disks
 echo "part / --fstype="xfs" --ondisk=$d1 --grow" >> /tmp/disks
 echo "part /boot --fstype="xfs" --ondisk=$d1 --size=1024" >> /tmp/disks

 %end
```

Bugzilla:1906870[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]

11.17. SUPPORTABILITY

The getattachment command fails to download multiple attachments at once

The **redhat-support-tool** command offers the **getattachment** subcommand for downloading attachments. However, **getattachment** is currently only able to download a single attachment and fails to download multiple attachments.

As a workaround, you can download multiple attachments one by one by passing the case number and UUID for each attachment in the **getattachment** subcommand.

Bugzilla:2064575

redhat-support-tool does not work with the FUTURE crypto policy

Because a cryptographic key used by a certificate on the Customer Portal API does not meet the requirements by the **FUTURE** system-wide cryptographic policy, the **redhat-support-tool** utility does not work with this policy level at the moment.

To work around this problem, use the **DEFAULT** crypto policy while connecting to the Customer Portal API.

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:2011413[1]

11.18. CONTAINERS

Running systemd within an older container image does not work

Running systemd within an older container image, for example, **centos:7**, does not work:

```
$ podman run --rm -ti centos:7 /usr/lib/systemd/systemd
 Storing signatures
 Failed to mount cgroup at /sys/fs/cgroup/systemd: Operation not permitted
 [!!!!!!] Failed to mount API filesystems, freezing.
```

To work around this problem, use the following commands:

```
# mkdir /sys/fs/cgroup/systemd
# mount none -t cgroup -o none,name=systemd /sys/fs/cgroup/systemd
# podman run --runtime /usr/bin/crun --annotation=run.oci.systemd.force_cgroup_v1=/sys/fs/cgroup -
-rm -ti centos:7 /usr/lib/systemd/systemd
```

Jira:RHELPLAN-96940[1]

CHAPTER 12. INTERNATIONALIZATION

12.1. RED HAT ENTERPRISE LINUX 8 INTERNATIONAL LANGUAGES

Red Hat Enterprise Linux 8 supports the installation of multiple languages and the changing of languages based on your requirements.

- East Asian Languages - Japanese, Korean, Simplified Chinese, and Traditional Chinese.

- European Languages – English, German, Spanish, French, Italian, Portuguese, and Russian.

The following table lists the fonts and input methods provided for various major languages.

Language	Default Font (Font Package)	Input Methods
English	dejavu-sans-fonts	
French	dejavu-sans-fonts	
German	dejavu-sans-fonts	
Italian	dejavu-sans-fonts	
Russian	dejavu-sans-fonts	
Spanish	dejavu-sans-fonts	
Portuguese	dejavu-sans-fonts	
Simplified Chinese	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-libpinyin, libpinyin
Traditional Chinese	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-libzhuyin, libzhuyin
Japanese	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-kkc, libkkc
Korean	google-noto-sans-cjk-ttc-fonts, google-noto-serif-cjk-ttc-fonts	ibus-hangul, libhangul

12.2. NOTABLE CHANGES TO INTERNATIONALIZATION IN RHEL 8

RHEL 8 introduces the following changes to internationalization compared to RHEL 7:

- Support for the **Unicode 11** computing industry standard has been added.

- Internationalization is distributed in multiple packages, which allows for smaller footprint installations. For more information, see Using langpacks.

- A number of **glibc** locales have been synchronized with Unicode Common Locale Data Repository (CLDR).

APPENDIX A. LIST OF TICKETS BY COMPONENT

Bugzilla and JIRA tickets are listed in this document for reference. The links lead to the release notes in this document that describe the tickets.

Component	Tickets
389-ds-base	Bugzilla:2188628, Bugzilla:2166332, Bugzilla:2166284, Bugzilla:2210491, Bugzilla:2220890, Bugzilla:2224505, Bugzilla:1817505
NetworkManager	Bugzilla:2144521, Bugzilla:2151987
Release Notes	Jira:RHELDOCS-16861, Jira:RHELDOCS-16755, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102
SLOF	Bugzilla:1910848
accel-config	Bugzilla:1843266
anaconda	Bugzilla:1770969, Bugzilla:1886985, Bugzilla:1656662, Bugzilla:2050140, Jira:RHEL-4707, Jira:RHEL-4711, Jira:RHEL-4744
ansible-freeipa	Bugzilla:2127901, Bugzilla:2175766, Bugzilla:2127906
apr	Bugzilla:1819607
audit	Bugzilla:2216666
authselect	Bugzilla:1892761
bacula	Jira:RHEL-6859
brltty	Bugzilla:2008197
clevis	Bugzilla:2209058
cloud-init	Bugzilla:2219528, Bugzilla:2230777, Jira:RHEL-12122
cockpit	Bugzilla:1666722
cockpit-appstream	Bugzilla:2212350, Bugzilla:2030836
cockpit-machines	Bugzilla:2173584
coreutils	Bugzilla:2030661
corosync-qdevice	Bugzilla:1784200

Component	Tickets
crash	Bugzilla:1906482
crash-ptdump-command	Bugzilla:1838927
createrepo_c	Bugzilla:1973588
crypto-policies	Bugzilla:2219912, Jira:RHEL-2345, Bugzilla:1919155, Bugzilla:1660839
cups-filters	Bugzilla:2118406
device-mapper-multipath	Bugzilla:2164871, Bugzilla:2022359, Bugzilla:2011699
distribution	Bugzilla:1657927
dnf	Bugzilla:2170093, Bugzilla:1986657
dnf-plugins-core	Bugzilla:2122587, Bugzilla:2092033
edk2	Bugzilla:1741615, Bugzilla:1935497
elfutils	Bugzilla:2182060, Bugzilla:2162495
fapolicyd	Jira:RHEL-628, Jira:RHEL-630, Jira:RHEL-829, Jira:RHEL-520, Bugzilla:2054741
fence-agents	Bugzilla:2187329, Bugzilla:1775847
firewalld	Bugzilla:1871860
fuse	Bugzilla:2171095
gcc	Bugzilla:2168205
gcc-toolset-12-gdb	Bugzilla:2172095
gcc-toolset-13	Bugzilla:2171898
gcc-toolset-13-annobin	Bugzilla:2171923, Bugzilla:2171921
gcc-toolset-13-binutils	Bugzilla:2171924
gcc-toolset-13-gcc	Bugzilla:2172091
gdb	Bugzilla:1853140

Component	Tickets
gfs2-utils	Bugzilla:2180782
glibc	Bugzilla:2180462
gnome-shell-extensions	Bugzilla:1717947
gnome-software	Bugzilla:1668760
gnutls	Bugzilla:2089817, Bugzilla:1628553
golang	Bugzilla:2185260
grafana	Bugzilla:2193250
grafana-pcp	Bugzilla:2193270
grub2	Bugzilla:1583445
grubby	Bugzilla:1900829
initscripts	Bugzilla:1875485
ipa	Bugzilla:2196425, Bugzilla:1821181, Jira:RHEL-4847, Jira:RHEL-4898, Jira:RHEL-12153, Bugzilla:1664719, Bugzilla:1664718, Bugzilla:2101770
ipmitool	Bugzilla:2224567, Jira:RHEL-6846
iproute	Jira:RHEL-424
kernel	Bugzilla:1989283, Bugzilla:2144529, Bugzilla:1753646, Bugzilla:2130727, Bugzilla:1868526, Bugzilla:1694705, Bugzilla:1730502, Bugzilla:1609288, Bugzilla:1602962, Bugzilla:1865745, Bugzilla:1906870, Bugzilla:1924016, Bugzilla:1942888, Bugzilla:1812577, Bugzilla:1910358, Bugzilla:1930576, Bugzilla:1793389, Bugzilla:1654962, Bugzilla:1940674, Bugzilla:1920086, Bugzilla:1971506, Bugzilla:2059262, Bugzilla:2050411, Bugzilla:2106341, Bugzilla:2189645, Bugzilla:1605216, Bugzilla:1519039, Bugzilla:1627455, Bugzilla:1501618, Bugzilla:1633143, Bugzilla:1814836, Bugzilla:1839311, Bugzilla:1696451, Bugzilla:1348508, Bugzilla:1837187, Bugzilla:1660337, Bugzilla:2041686, Bugzilla:1836977, Bugzilla:1878207, Bugzilla:1665295, Bugzilla:1871863, Bugzilla:1569610, Bugzilla:1794513
kernel / Networking / IPSec	Jira:RHEL-1257
kernel / Networking / NIC Drivers	Jira:RHEL-11398

Component	Tickets
kernel / **Virtualization** / **KVM**	Jira:RHEL-2451
kernel-rt / **Other**	Jira:RHEL-9318
kexec-tools	Bugzilla:2173791, Bugzilla:2111855
kmod	Bugzilla:2103605
krb5	Bugzilla:2211390, Jira:RHEL-4910, Bugzilla:2125318, Bugzilla:1877991
leapp-repository	Bugzilla:2097003
libdnf	Bugzilla:2155713
libgnome-keyring	Bugzilla:1607766
libguestfs	Bugzilla:1554735
libnftnl	Bugzilla:2211096
libpfm	Bugzilla:2185653
libreswan	Bugzilla:1989050
libselinux-python-2.8-module	Bugzilla:1666328
libvirt	Bugzilla:1664592, Bugzilla:1332758, Bugzilla:2143160, Bugzilla:1528684
llvm-toolset	Bugzilla:2178806
lvm2	Bugzilla:1496229, Bugzilla:1768536
mariadb	Bugzilla:1942330
mesa	Bugzilla:1886147
nfs-utils	Bugzilla:2081114, Bugzilla:1592011
nftables	Bugzilla:2061942
nodejs	Bugzilla:2186718
nss	Bugzilla:1817533, Bugzilla:1645153

Component	Tickets
nss_nis	Bugzilla:1803161
opencryptoki	Bugzilla:2159697
opencv	Bugzilla:1886310
openmpi	Bugzilla:1866402
opensc	Bugzilla:2097048, Jira:RHEL-4077, Bugzilla:1947025
openscap	Bugzilla:2217441, Bugzilla:2161499
openssh	Bugzilla:2044354
openssl	Bugzilla:1810911
osbuild-composer	Jira:RHEL-4649
oscap-anaconda-addon	Jira:RHEL-1826, Bugzilla:1843932, Bugzilla:1665082, Jira:RHEL-1810
pacemaker	Bugzilla:1876173, Bugzilla:2160206, Bugzilla:2078611, Bugzilla:2030869, Bugzilla:2010084, Bugzilla:1632951, Bugzilla:1578820, Bugzilla:1931023, Bugzilla:2168633
papi	Bugzilla:2111982, Bugzilla:2161146
pcs	Bugzilla:2166294, Bugzilla:2179010, Bugzilla:2189958, Bugzilla:2166289, Bugzilla:1619620, Bugzilla:1851335
perl-HTTP-Tiny	Bugzilla:2228409
pki-core	Bugzilla:1729215, Jira:RHEL-13125, Bugzilla:1628987
podman	Jira:RHELPLAN-154313, Jira:RHELPLAN-154431, Jira:RHELPLAN-154440, Jira:RHELPLAN-154443, Jira:RHELPLAN-163002, Jira:RHELPLAN-160659, Jira:RHELPLAN-154428
postfix	Bugzilla:1787010, Bugzilla:1711885
pykickstart	Bugzilla:1637872
python3.11-lxml	Bugzilla:2157673
python36-3.6-module	Bugzilla:2165702

Component	Tickets
qemu-kvm	Jira:RHEL-13336, Bugzilla:1740002, Bugzilla:1719687, Bugzilla:1966475, Bugzilla:1792683, Bugzilla:2177957, Bugzilla:1651994
rear	Bugzilla:2233526, Bugzilla:1925531, Bugzilla:2083301
redhat-support-tool	Bugzilla:2064575
resource-agents	Bugzilla:2040110, Bugzilla:2049319, Bugzilla:2039692, Bugzilla:2181019
restore	Bugzilla:1997366
rhel-system-roles	Bugzilla:2151371, Bugzilla:2224387, Bugzilla:2190483, Bugzilla:2141961, Bugzilla:2181661, Bugzilla:2211272, Bugzilla:2211723, Bugzilla:2211778, Bugzilla:2216759, Bugzilla:2218204, Bugzilla:2224388, Bugzilla:2218595, Bugzilla:2211273, Bugzilla:2140880, Bugzilla:2192343, Bugzilla:2222809, Jira:RHEL-866, Jira:RHEL-858, Bugzilla:2168738, Bugzilla:2186057, Bugzilla:2209441, Bugzilla:2216521, Bugzilla:2224094, Bugzilla:2224648, Bugzilla:2226077, Bugzilla:2193057, Bugzilla:2222433, Bugzilla:2232391, Bugzilla:2232392, Jira:RHEL-899, Jira:RHEL-907, Jira:RHEL-918, Jira:RHEL-1398, Jira:RHEL-1496, Jira:RHEL-1500, Bugzilla:2186908, Bugzilla:2021685, Bugzilla:2006081
rpm	Bugzilla:1688849
rsyslog	Jira:RHELPLAN-160541, Bugzilla:1679512, Jira:RHELPLAN-10431
rust-toolset	Bugzilla:2191740, Bugzilla:2213875
samba	Bugzilla:2190417, Bugzilla:2009213, Jira:RHELPLAN-13195
scap-security-guide	Bugzilla:2155789, Bugzilla:2157877, Bugzilla:2167999, Bugzilla:2221695, Bugzilla:2129100, Bugzilla:2169857, Bugzilla:2130185, Bugzilla:2175684, Bugzilla:2175882, Bugzilla:2184487, Bugzilla:2192893, Bugzilla:2170530, Bugzilla:2176008, Bugzilla:2209073, Bugzilla:2222583, Bugzilla:2028428, Bugzilla:2118758, Jira:RHEL-1804, Jira:RHEL-1897
selinux-policy	Bugzilla:2172541, Bugzilla:2184348, Bugzilla:2196524, Bugzilla:2166153, Bugzilla:1461914
sos	Bugzilla:2011413
spice	Bugzilla:1849563
sssd	Bugzilla:2065692, Bugzilla:2056483, Bugzilla:1947671
subscription-manager	Bugzilla:2170082

Component	Tickets
sysstat	Jira:RHEL-12008
systemtap	Bugzilla:2186932, Bugzilla:2126805
tang	Bugzilla:2188743
tuned	Bugzilla:2113900
udica	Bugzilla:1763210
udisks2	Bugzilla:2213193
valgrind	Bugzilla:2124345
vdo	Bugzilla:1949163
virt-manager	Bugzilla:2026985
vsftpd	Bugzilla:2069733
wayland	Bugzilla:1673073
webkit2gtk3	Jira:RHEL-4158
which	Bugzilla:2140566
xorg-x11-server	Bugzilla:1698565

Component	Tickets
other	Jira:RHELDOCS-16405, Bugzilla:2232558, Jira:RHELDOCS-16247, Jira:RHELDOCS-16474, Jira:RHELDOCS-16462, Jira:RHELPLAN-156196, Jira:RHELDOCS-16339, Jira:RHELDOCS-16367, Jira:RHELDOCS-17369, Bugzilla:2236183, Bugzilla:2025814, Bugzilla:2077770, Bugzilla:1777138, Bugzilla:1640697, Bugzilla:1697896, Bugzilla:1961722, Bugzilla:1659609, Bugzilla:1687900, Bugzilla:1757877, Bugzilla:1741436, Jira:RHELPLAN-27987, Jira:RHELPLAN-34199, Jira:RHELPLAN-57914, Jira:RHELPLAN-96940, Bugzilla:1974622, Bugzilla:2028361, Bugzilla:2041997, Bugzilla:2035158, Jira:RHELPLAN-109613, Bugzilla:2126777, Bugzilla:1690207, Bugzilla:1559616, Bugzilla:1889737, Bugzilla:1906489, Bugzilla:1769727, Jira:RHELPLAN-27394, Jira:RHELPLAN-27737, Jira:RHELDOCS-16861, Bugzilla:1642765, Bugzilla:1646541, Bugzilla:1647725, Jira:RHELDOCS-17380, Bugzilla:1932222, Bugzilla:1686057, Bugzilla:1748980, Jira:RHELPLAN-71200, Jira:RHELPLAN-45858, Bugzilla:1871025, Bugzilla:1871953, Bugzilla:1874892, Bugzilla:1916296, Jira:RHELPLAN-100400, Bugzilla:1926114, Bugzilla:1904251, Bugzilla:2011208, Jira:RHELPLAN-59825, Bugzilla:1920624, Jira:RHELPLAN-70700, Bugzilla:1929173, Jira:RHELPLAN-85066, Jira:RHELPLAN-98983, Bugzilla:2009113, Bugzilla:1958250, Bugzilla:2038929, Bugzilla:2006665, Bugzilla:2029338, Bugzilla:2061288, Bugzilla:2060759, Bugzilla:2055826, Bugzilla:2059626, Jira:RHELPLAN-133171, Bugzilla:2142499, Jira:RHELDOCS-16755, Jira:RHELPLAN-146398, Jira:RHELPLAN-153267, Bugzilla:2225332, Jira:RHELPLAN-147538, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102, Jira:RHELDOCS-16300

APPENDIX B. REVISION HISTORY

0.0-11

Fri May 10 2024, Brian Angelica (bangelic@redhat.com)

- Updated Tech Preview in BZ#1690207.

0.0-10

Thu May 9 2024, Gabriela Fialova (gfialova@redhat.com)

- Updated a known issue BZ#1730502 (Storage).

0.0-9

Mon April 29 2024, Gabriela Fialova (gfialova@redhat.com)

- Added an enhancement BZ#2093355 (Security).

0.0-8

Mon March 4 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a bug fix Jira:SSSD-6096 (Identity Management).

0.0-7

Thu February 29 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a deprecated functionality Jira:RHELDOCS-17641 (Networking).

0.0-6

Tue February 13 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a deprecated functionality Jira:RHELDOCS-17573 (Identity Management).

0.0-5

Fri February 2 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a known issue BZ#1834716 (Security).

- Updated the Jira:RHELDOCS-16755 deprecated functionality note (Containers).

0.0-4

Fri January 19 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added an enhancement related to Python Jira:RHELDOCS-17369 (Dynamic programming languages, web and database servers).

- Added an enhancement Jira:RHELDOCS-16367 (The web console).

0.0-3

Wed January 10 2024, Lucie Vařáková (lvarakova@redhat.com)

- Added a rebase BZ#2196425 (Identity Management).

- Updated the Jira:RHELPLAN-156196 new feature description (Supportability).

- Added deprecated functionality Jira:RHELDOCS-17380 (Security).

- Other minor updates.

0.0-2

Thu November 16 2023, Lenka Špačková (lspackova@redhat.com)

- **Node.js 20** is now fully supported (BZ#2186718).

0.0-1

Wed November 15 2023, Lucie Vařáková (lvarakova@redhat.com)

- Release of the Red Hat Enterprise Linux 8.9 Release Notes.

0.0-0

Wed September 27 2023, Lucie Vařáková (lvarakova@redhat.com)

- Release of the Red Hat Enterprise Linux 8.9 Beta Release Notes.