



Red Hat Enterprise Linux 8.10

8.10 Release Notes

Release Notes for Red Hat Enterprise Linux 8.10

Red Hat Enterprise Linux 8.10 8.10 Release Notes

Release Notes for Red Hat Enterprise Linux 8.10

Legal Notice

Copyright © 2024 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 8.10 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	5
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	6
CHAPTER 1. OVERVIEW	7
1.1. MAJOR CHANGES IN RHEL 8.10	7
Installer and image creation	7
Security	7
Dynamic programming languages, web and database servers	7
Identity Management	7
Containers	8
1.2. IN-PLACE UPGRADE AND OS CONVERSION	8
In-place upgrade from RHEL 7 to RHEL 8	8
In-place upgrade from RHEL 6 to RHEL 8	9
In-place upgrade from RHEL 8 to RHEL 9	9
Conversion from a different Linux distribution to RHEL	9
1.3. RED HAT CUSTOMER PORTAL LABS	9
1.4. ADDITIONAL RESOURCES	10
CHAPTER 2. ARCHITECTURES	11
CHAPTER 3. DISTRIBUTION OF CONTENT IN RHEL 8	12
3.1. INSTALLATION	12
3.2. REPOSITORIES	12
3.3. APPLICATION STREAMS	13
3.4. PACKAGE MANAGEMENT WITH YUM/DNF	13
CHAPTER 4. NEW FEATURES	14
4.1. INSTALLER AND IMAGE CREATION	14
4.2. SECURITY	14
4.3. SHELLS AND COMMAND-LINE TOOLS	17
4.4. INFRASTRUCTURE SERVICES	17
4.5. NETWORKING	18
4.6. KERNEL	18
4.7. BOOT LOADER	20
4.8. FILE SYSTEMS AND STORAGE	20
4.9. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	21
4.10. COMPILERS AND DEVELOPMENT TOOLS	27
4.11. IDENTITY MANAGEMENT	30
4.12. THE WEB CONSOLE	33
4.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES	34
4.14. VIRTUALIZATION	38
4.15. RHEL IN CLOUD ENVIRONMENTS	38
4.16. CONTAINERS	39
CHAPTER 5. AVAILABLE BPF FEATURES	42
CHAPTER 6. BUG FIXES	56
6.1. INSTALLER AND IMAGE CREATION	56
6.2. SECURITY	56
6.3. SOFTWARE MANAGEMENT	57
6.4. SHELLS AND COMMAND-LINE TOOLS	58
6.5. KERNEL	59

6.6. FILE SYSTEMS AND STORAGE	60
6.7. HIGH AVAILABILITY AND CLUSTERS	60
6.8. COMPILERS AND DEVELOPMENT TOOLS	61
6.9. IDENTITY MANAGEMENT	62
6.10. RED HAT ENTERPRISE LINUX SYSTEM ROLES	65
6.11. VIRTUALIZATION	67
CHAPTER 7. TECHNOLOGY PREVIEWS	68
7.1. INFRASTRUCTURE SERVICES	68
7.2. NETWORKING	68
7.3. KERNEL	69
7.4. FILE SYSTEMS AND STORAGE	70
7.5. HIGH AVAILABILITY AND CLUSTERS	72
7.6. IDENTITY MANAGEMENT	73
7.7. DESKTOP	75
7.8. GRAPHICS INFRASTRUCTURES	76
7.9. VIRTUALIZATION	76
7.10. RHEL IN CLOUD ENVIRONMENTS	78
7.11. CONTAINERS	78
CHAPTER 8. DEPRECATED FUNCTIONALITY	80
8.1. INSTALLER AND IMAGE CREATION	80
8.2. SECURITY	81
8.3. SUBSCRIPTION MANAGEMENT	83
8.4. SOFTWARE MANAGEMENT	83
8.5. SHELLS AND COMMAND-LINE TOOLS	83
8.6. INFRASTRUCTURE SERVICES	85
8.7. NETWORKING	85
8.8. KERNEL	86
8.9. BOOT LOADER	87
8.10. FILE SYSTEMS AND STORAGE	87
8.11. HIGH AVAILABILITY AND CLUSTERS	89
8.12. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	89
8.13. COMPILERS AND DEVELOPMENT TOOLS	89
8.14. IDENTITY MANAGEMENT	90
8.15. DESKTOP	93
8.16. GRAPHICS INFRASTRUCTURES	94
8.17. THE WEB CONSOLE	94
8.18. RED HAT ENTERPRISE LINUX SYSTEM ROLES	95
8.19. VIRTUALIZATION	96
8.20. CONTAINERS	97
8.21. DEPRECATED PACKAGES	99
8.22. DEPRECATED AND UNMAINTAINED DEVICES	140
CHAPTER 9. KNOWN ISSUES	145
9.1. INSTALLER AND IMAGE CREATION	145
9.2. SECURITY	147
9.3. RHEL FOR EDGE	154
9.4. SUBSCRIPTION MANAGEMENT	154
9.5. SOFTWARE MANAGEMENT	154
9.6. SHELLS AND COMMAND-LINE TOOLS	155
9.7. INFRASTRUCTURE SERVICES	157
9.8. NETWORKING	158
9.9. KERNEL	158

9.10. FILE SYSTEMS AND STORAGE	164
9.11. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS	166
9.12. IDENTITY MANAGEMENT	167
9.13. DESKTOP	170
9.14. GRAPHICS INFRASTRUCTURES	170
9.15. RED HAT ENTERPRISE LINUX SYSTEM ROLES	172
9.16. VIRTUALIZATION	173
9.17. RHEL IN CLOUD ENVIRONMENTS	177
9.18. SUPPORTABILITY	179
9.19. CONTAINERS	180
CHAPTER 10. INTERNATIONALIZATION	181
10.1. RED HAT ENTERPRISE LINUX 8 INTERNATIONAL LANGUAGES	181
10.2. NOTABLE CHANGES TO INTERNATIONALIZATION IN RHEL 8	181
APPENDIX A. LIST OF TICKETS BY COMPONENT	183
APPENDIX B. REVISION HISTORY	191

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 8.10

Installer and image creation

Key highlights for RHEL image builder:

- You can create different partitioning modes, such as **auto-lvm**, **lvm**, and **raw**.
- You can customize tailoring options for a profile and add it to your blueprint customizations by using selected and unselected options, to add and remove rules.

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

Security

SCAP Security Guide 0.1.72 contains updated CIS profiles, a profile aligned with the PCI DSS policy version 4.0, and profiles for the latest DISA STIG policies.

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

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

The **OpenSSL** TLS toolkit now contains API-level protections against Bleichenbacher-like attacks on the RSA PKCS #1 v1.5 decryption process.

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

Dynamic programming languages, web and database servers

Later versions of the following Application Streams are now available:

- Python 3.12
- Ruby 3.3
- PHP 8.2
- nginx 1.24
- MariaDB 10.11
- PostgreSQL 16

The following components have been upgraded:

- **Git** to version 2.43.0
- **Git LFS** to version 3.4.1

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

Identity Management

Identity Management (IdM) in RHEL 8.10 introduces delegating user authentication to external identity providers (IdPs) that support the OAuth 2 Device Authorization Grant flow. This is now a fully supported feature.

After performing authentication and authorization at the external IdP, the IdM user receives a Kerberos ticket with single sign-on capabilities.

For more information, see [New Features - Identity Management](#)

Containers

Notable changes include:

- The **podman farm build** command for creating multi-architecture container images is available as a Technology Preview.
- Podman now supports **containers.conf** modules to load a predetermined set of configurations.
- The Container Tools packages have been updated.
- Podman v4.9 RESTful API now displays data of progress when you pull or push an image to the registry.
- SQLite is now fully supported as a default database backend for Podman.
- **Containerfile** now supports multi-line HereDoc instructions.
- **pasta** as a network name has been deprecated.
- The BoltDB database backend has been deprecated.
- The **container-tools:4.0** module has been deprecated.
- The Container Network Interface (CNI) network stack is deprecated and will be removed in a future release.

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

1.2. IN-PLACE UPGRADE 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 8.8 and RHEL 8.10 on the 64-bit Intel, IBM POWER 8 (little endian), and IBM Z architectures
- From RHEL 7.9 to RHEL 8.8 and RHEL 8.10 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](#).

For information regarding how Red Hat supports the in-place upgrade process, see the [In-place upgrade Support Policy](#).

Notable enhancements include:

- New logic has been implemented to determine the expected states of the **systemd** services after the upgrade.
- Locally stored DNF repositories can now be used for the in-place upgrade.
- You can now configure DNF to be able to upgrade by using proxy.
- Issues with performing the in-place upgrade with custom DNF repositories accessed by using HTTPS have been fixed.
- If the **/etc/pki/tls/openssl.cnf** configuration file has been modified, the file is now replaced with the target default OpenSSL configuration file during the upgrade to prevent issues after the upgrade. See the pre-upgrade report for more information.

In-place upgrade from RHEL 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 CentOS Linux 7 or Oracle Linux 7, you can convert your operating system to RHEL and then perform an in-place upgrade to RHEL 8.

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 Ops 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 8.10 is distributed with the kernel version 4.18.0-553, 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 8.10.

4.1. INSTALLER AND IMAGE CREATION

Ability to use partitioning mode on the blueprint filesystem customization

With this update, while using RHEL image builder, you can customize your blueprint with the chosen filesystem customization. You can choose one of the following partition modes while you create an image:

- Default: **auto-lvm**
- LVM: the image uses Logical Volume Manager (LVM) even without extra partitions
- Raw: the image uses raw partitioning even with extra partitions

Jira:RHELDPCS-16337^[1]

Filesystem customization policy changes in image builder

The following policy changes are in place when using the RHEL image builder filesystem customization in blueprints:

Currently, **mountpoint** and minimum partition **minsize** can be set. The following image types do not support filesystem customizations: **image-installeredge-installeredge-simplified-installer** The following image types do not create partitioned operating systems images. Customizing their filesystem is meaningless: **edge-commitedge-containertarcontainer** The blueprint now supports the **mountpoint** customization for **tpm** and its sub-directories.

Jira:RHELDPCS-17261^[1]

4.2. SECURITY

SCAP Security Guide rebased to 0.1.72

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

- CIS profiles are updated to align with the latest benchmarks.
- The PCI DSS profile is aligned with the PCI DSS policy version 4.0.
- STIG profiles are aligned with the latest DISA STIG policies.

For additional information, see the [SCAP Security Guide release notes](#).

Jira:RHEL-25250^[1]

OpenSSL now contains protections against Bleichenbacher-like attacks

This release of the OpenSSL TLS toolkit introduces API-level protections against Bleichenbacher-like attacks on the RSA PKCS #1 v1.5 decryption process. The RSA decryption now returns a randomly generated deterministic message instead of an error if it detects an error when checking padding during


a PKCS #1 v1.5 decryption. The change provides general protection against vulnerabilities such as [CVE-2020-25659](#) and [CVE-2020-25657](#).

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

Jira:RHEL-17689^[1]

librdkafka rebased to 1.6.1

The **librdkafka** implementation of the Apache Kafka protocol has been rebased to upstream version 1.6.1. This is the first major feature release for RHEL 8. The rebase provides many important enhancements and bug fixes. For all relevant changes, see the **CHANGELOG.md** document provided in the **librdkafka** package.



NOTE

This update changes configuration defaults and deprecates some configuration properties. Read the Upgrade considerations section in **CHANGELOG.md** for more details. The API (C & C++) and ABI © in this version are compatible with older versions of **librdkafka**, but some changes to the configuration properties may require changes to existing applications.

Jira:RHEL-12892^[1]

libkcapi rebased to 1.4.0

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

- Added the **sm3sum** and **sm3hmac** tools.
- Added the **kcapi_md_sm3** and **kcapi_md_hmac_sm3** APIs.
- Added SM4 convenience functions.
- Fixed support for link-time optimization (LTO).
- Fixed LTO regression testing.
- Fixed support for AEAD encryption of an arbitrary size with **kcapi-enc**.

Jira:RHEL-5366^[1]

stunnel rebased to 5.71

The **stunnel** TLS/SSL tunneling service has been rebased to upstream version 5.71. This update changes the behavior of OpenSSL 1.1 and later versions in FIPS mode. If OpenSSL is in FIPS mode and **stunnel** default FIPS configuration is set to **no**, **stunnel** adapts to OpenSSL and FIPS mode is enabled.

Additional new features include:

- Added support for modern PostgreSQL clients.
- You can use the **protocolHeader** service-level option to insert custom **connect** protocol negotiation headers.

- You can use the **protocolHost** option to control the client SMTP protocol negotiation HELO/EHLO value.
- Added client-side support for Client-side **protocol = ldap**.
- You can now configure session resumption by using the service-level **sessionResume** option.
- Added support to request client certificates in server mode with **CPath** (previously, only **CFile** was supported).
- Improved file reading and logging performance.
- Added support for configurable delay for the **retry** option.
- In client mode, OCSP stapling is requested and verified when **verifyChain** is set.
- In server mode, OCSP stapling is always available.
- Inconclusive OCSP verification breaks TLS negotiation. You can disable this by setting **OCSPRequire = no**.

[Jira:RHEL-2340^{\[1\]}](#)

OpenSSH limits artificial delays in authentication

OpenSSH's response after login failure is artificially delayed to prevent user enumeration attacks. This update introduces an upper limit so that such artificial delays do not become excessively long when remote authentication takes too long, for example in privilege access management (PAM) processing.

[Jira:RHEL-1684](#)

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

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

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

[Jira:RHEL-15300^{\[1\]}](#)

audit rebased to 3.1.2

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

- The **ausearch** library now interprets unnamed and anonymous sockets.
- You can use the new keyword **this-hour** in the **start** and **end** options of the **ausearch** and **ausearch** tools.
- User-friendly keywords for signals have been added to the **auditctl** program.
- Handling of corrupt logs in **ausearch** has been improved.
- The **ProtectControlGroups** option is now disabled by default in the **auditd** service.
- Rule checking for the exclude filter has been fixed.

- The interpretation of **OPENAT2** fields has been enhanced.
- The **audispd af_unix** plugin has been moved to a standalone program.
- The Python binding has been changed to prevent setting Audit rules from the Python API. This change was made due to a bug in the Simplified Wrapper and Interface Generator (SWIG).

Jira:RHEL-15001^[1]

4.3. SHELLS AND COMMAND-LINE TOOLS

openCryptoki rebased to version 3.22.0

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

- Added support for the **AES-XTS** key type by using the **CPACF** protected keys.
- Added support for managing certificate objects.
- Added support for public sessions with the **no-login** option.
- Added support for logging in as the Security Officer (SO).
- Added support for importing and exporting the **Edwards** and **Montgomery** keys.
- Added support for importing the **RSA-PSS** keys and certificates.
- For security reasons, the 2 key parts of an AES-XTS key should not be the same. This update adds checks to the key generation and import process to ensure this.
- Various bug fixes have been implemented.

Jira:RHEL-11413^[1]

4.4. INFRASTRUCTURE SERVICES

chrony rebased to version 4.5

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

- Added periodic refresh of IP addresses of Network Time Protocol (NTP) sources specified by hostname. The default interval is two weeks and it can be disabled by adding **refresh 0** to the **chrony.conf** file.
- Improved automatic replacement of unreachable NTP sources.
- Improved logging of important changes made by the **chronyc** utility.
- Improved logging of source selection failures and falsetickers.
- Added the **hwtsttimeout** directive to configure timeout for late hardware transmit timestamps.
- Added experimental support for corrections provided by Precision Time Protocol (PTP) transparent clocks to reach accuracy of PTP with hardware timestamping.
- Fixed the **presend** option in **interleaved** mode.

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

[Jira:RHEL-21069](#)

linuxptp rebased to version 4.2

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

- Added support for multiple domains in the **phc2sys** utility.
- Added support for notifications on clock updates and changes in the Precision Time Protocol (PTP) parent dataset, for example, clock class.
- Added support for PTP Power Profile, namely IEEE C37.238-2011 and IEEE C37.238-2017.

[Jira:RHEL-21326^{\[1\]}](#)

4.5. NETWORKING

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

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

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

[Jira:RHEL-6113^{\[1\]}](#)

nispor rebased to version 1.2.10

The **nispor** packages have been upgraded to upstream version 1.2.10, which provides a number of enhancements and bug fixes over the previous version:

- Added support for **NetStateFilter** to use the kernel filter on network routes and interfaces.
- Single Root Input and Output Virtualization (SR-IOV) interfaces can query SR-IOV Virtual Function (SR-IOV VF) information per (VF).
- Newly supported bonding options: **lACP_active**, **arp_missed_max**, and **ns_ip6_target**.

[Bugzilla:2153166](#)

4.6. KERNEL

Kernel version in RHEL 8.10

Red Hat Enterprise Linux 8.10 is distributed with the kernel version 4.18.0-553.

rtla rebased to version 6.6 of the upstream kernel source code

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

- Added the **-C** option to specify additional control groups for **rtla** threads to run in, apart from the main **rtla** thread.
- Added the **--house-keeping** option to place **rtla** threads on a housekeeping CPU and to put measurement threads on different CPUs.
- Added support to the **timerlat** tracer so that you can run **timerlat hist** and **timerlat top** threads in user space.

Jira:RHEL-10081^[1]

rteval was upgraded to the upstream version 3.7

With this update, the **rteval** utility has been upgraded to the upstream version 3.7. The most significant feature in this update concerns the **isolcpus** kernel parameter. This includes the ability to detect and use the **isolcpus** mechanism for measurement modules in **rteval**. As a result, it is easier for **isolcpus** users to use **rteval** to get accurate latency numbers and to achieve best latency results measured on a realtime kernel.

Jira:RHEL-8967^[1]

SGX is now fully supported

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

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

- Modifying EPCM permissions of regular enclave pages that belong to an initialized enclave.
- Dynamic addition of regular enclave pages to an initialized enclave.
- Expanding an initialized enclave to accommodate more threads.
- Removing regular and TCS pages from an initialized enclave.

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

Bugzilla:2041881^[1]

The Intel data streaming accelerator driver is now fully supported

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

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

Jira:RHEL-10097^[1]

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

With the **rteval** utility, you can add (using the + sign) or subtract (using the - sign) CPUs to the default measurement CPU list when using the **--measurement-cpulists** parameter, instead of having to specify an entire new list. Additionally, **--measurement-run-on-isolcpus** is introduced for adding the set of all

isolated CPUs to the default measurement CPU list. This options covers the most common usecase of a real-time application running on isolated CPUs. Other usecases require a more generic feature. For example, some real-time applications used one isolated CPU for housekeeping, requiring it to be excluded from the default measurement CPU list. As a result, you can now not only add, but also remove arbitrary CPUs from the default measurement CPU list in a flexible way. Removing takes precedence over adding. This rule applies to both, CPUs specified with +/- signs and to those defined with **--measurement-run-on-isolcpus**.

Jira:RHEL-21926^[1]

4.7. BOOT LOADER

DEP/NX support in the pre-boot stage

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

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

Jira:RHEL-15856^[1]

Support for TD RTMR measurement in GRUB and shim

Intel® Trust Domain Extension (Intel® TDX) is a confidential computing technology that deploys hardware-isolated virtual machines (VMs) called Trust Domains (TDs).

TDX extends the Virtual Machine Extensions (VMX) instructions and the Multi-key Total Memory Encryption (MKTME) feature with the TD VM guest. In a TD guest VM, all components in the boot chain, such as **grub2** and **shim**, must log the event and measurement hash to runtime measurement registers (RTMR).

TD guest runtime measurement in RTMR is the base for attestation applications. Applications on the TD guest rely on TD measurement to provide trust evidence to get confidential information, such as the key from the relaying part through the attestation service.

With this release, the GRUB and **shim** boot loaders now support the TD measurement protocol.

For more information about Intel® TDX, see [Documentation for Intel® Trust Domain Extensions](#).

Jira:RHEL-15583^[1]

4.8. FILE SYSTEMS AND STORAGE

The Storage RHEL System Roles now support shared LVM device management

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

[Jira:RHEL-14022](#)

multipathd now supports detecting FPIN-Li events for NVMe devices

Previously, the **multipathd** command would only monitor Integrity Fabric Performance Impact

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

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

[Jira:RHEL-6677](#)

4.9. DYNAMIC PROGRAMMING LANGUAGES, WEB AND DATABASE SERVERS

Python 3.12 available in RHEL 8

RHEL 8.10 introduces Python 3.12, provided by the new package **python3.12** and a suite of packages built for it, as well as the **ubi8/python-312** container image.

Notable enhancements compared to the previously released Python 3.11 include:

- Python introduces a new **type** statement and new type parameter syntax for generic classes and functions.
- Formatted string literal (f-strings) have been formalized in the grammar and can now be integrated into the parser directly.
- Python now provides a unique per-interpreter global interpreter lock (GIL).
- You can now use the buffer protocol from Python code.
- To improve security, the builtin **hashlib** implementations of the SHA1, SHA3, SHA2-384, SHA2-512, and MD5 cryptographic algorithms have been replaced with formally verified code from the HACLS* project. The builtin implementations remain available as fallback if OpenSSL does not provide them.
- Dictionary, list, and set comprehensions in **CPython** are now inlined. This significantly increases the speed of a comprehension execution.
- **CPython** now supports the Linux **perf** profiler.
- **CPython** now provides stack overflow protection on supported platforms.

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

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

To run the interpreter, use, for example:

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

See [Installing and using Python](#) for more information.

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

[Jira:RHEL-14942](#)

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

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

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

```
export PYTHON_EMAIL_DISABLE_STRICT_ADDR_PARSING=true
```

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

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

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

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

[Jira:RHELDOS-17369^{\[1\]}](#)

A new module stream: ruby:3.3

RHEL 8.10 introduces Ruby 3.3.0 in a new **ruby:3.3** module stream. This version provides a number of performance improvements, bug and security fixes, and new features over **Ruby 3.1** distributed with RHEL 8.7.

Notable enhancements include:

- You can use the new **Prism** parser instead of **Ripper**. **Prism** is a portable, error tolerant, and maintainable recursive descent parser for the Ruby language.
- YJIT, the Ruby just-in-time (JIT) compiler implementation, is no longer experimental and it provides major performance improvements.
- The **Regexp** matching algorithm has been improved to reduce the impact of potential Regular Expression Denial of Service (ReDoS) vulnerabilities.
- The new experimental RJIT (a pure-Ruby JIT) compiler replaces MJIT. Use YJIT in production.
- A new M:N thread scheduler is now available.

Other notable changes:

- You must now use the **Lrama** LALR parser generator instead of **Bison**.
- Several deprecated methods and constants have been removed.
- The **Racc** gem has been promoted from a default gem to a bundled gem.

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

```
# yum module install ruby:3.3
```

If you want to upgrade from an earlier **ruby** module stream, see [Switching to a later stream](#).

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

Jira:RHEL-17090^[1]

A new module stream: **php:8.2**

RHEL 8.10 adds PHP 8.2, which provides a number of bug fixes and enhancements over version 8.0.

With **PHP 8.2**, you can:

- Define a custom type that is limited to one of a discrete number of possible values using the Enumerations (Enums) feature.
- Declare a property with the **readonly** modifier to prevent modification of the property after initialization.
- Use fibers, full-stack, and interruptible functions.
- Use readonly classes.
- Declare several new standalone types.
- Use a new **Random** extension.
- Define constraints in traits.

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

```
# yum module install php:8.2
```

If you want to upgrade from an earlier **php** stream, see [Switching to a later stream](#).

For details regarding PHP usage on RHEL 8, see [Using the PHP scripting language](#).

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

Jira:RHEL-14705^[1]

The **name()** method of the **perl-DateTime-TimeZone** module now returns the time zone name

The **perl-DateTime-TimeZone** module has been updated to version 2.62, which changed the value that is returned by the **name()** method from the time zone alias to the main time zone name.

For more information and an example, see the Knowledgebase article [Change in the perl-DateTime-TimeZone API related to time zone name and alias](#).

Jira:RHEL-35685

A new module stream: **nginx:1.24**

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

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

- Encryption keys are now automatically rotated for TLS session tickets when using shared memory in the **ssl_session_cache** directive.
- Memory usage has been optimized in configurations with Secure Sockets Layer (SSL) proxy.
- You can now disable looking up IPv4 addresses while resolving by using the **ipv4=off** parameter of the **resolver** directive.
- nginx now supports the **\$proxy_protocol_tlv_*** variables, which store the values of the Type-Length-Value (TLV) fields that appear in the PROXY v2 TLV protocol.
- The **ngx_http_gzip_static_module** module now supports byte ranges.

Other changes:

- Header lines are now represented as linked lists in the internal API.
- nginx now concatenates identically named header strings passed to the FastCGI, SCGI, and uwsgi back ends in the **\$r->header_in()** method of the **ngx_http_perl_module**, and during lookups of the **\$http_...**, **\$sent_http_...**, **\$sent_trailer_...**, **\$upstream_http_...**, and **\$upstream_trailer_...** variables.
- nginx now displays a warning if protocol parameters of a listening socket are redefined.
- nginx now closes connections with lingering if pipelining was used by the client.
- The logging level of various SSL errors has been lowered, for example, from **Critical** to **Informational**.

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

```
# yum module install nginx:1.24
```

To upgrade from an earlier **nginx** stream, [switch to a later stream](#).

For more information, see [Setting up and configuring NGINX](#).

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

Jira:RHEL-14714^[1]

A new module stream: **mariadb:10.11**

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

- A new **sys_schema** feature.
- Atomic Data Definition Language (DDL) statements.

- A new **GRANT ... TO PUBLIC** privilege.
- Separate **SUPER** and **READ ONLY ADMIN** privileges.
- A new **UUID** database data type.
- Support for the Secure Socket Layer (SSL) protocol version 3; the MariaDB server now requires correctly configured SSL to start.
- Support for the natural sort order through the **natural_sort_key()** function.
- A new **SFORMAT** function for arbitrary text formatting.
- Changes to the UTF-8 charset and the UCA-14 collation.
- **systemd** socket activation files available in the `/usr/share/` directory. Note that they are not a part of the default configuration in RHEL as opposed to upstream.
- Error messages containing the **MariaDB** string instead of **MySQL**.
- Error messages available in the Chinese language.
- Changes to the default logrotate file.
- For MariaDB and MySQL clients, the connection property specified on the command line (for example, **--port=3306**), now forces the protocol type of communication between the client and the server, such as **tcp**, **socket**, **pipe**, or **memory**.

For more information about changes in MariaDB 10.11, see [Notable differences between MariaDB 10.5 and MariaDB 10.11](#).

For more information about MariaDB, see [Using MariaDB](#).

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

```
# yum module install mariadb:10.11
```

If you want to upgrade from the **mariadb:10.5** module stream, see [Upgrading from MariaDB 10.5 to MariaDB 10.11](#).

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

[Jira:RHEL-3637](#)

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

RHEL 8.10 introduces PostgreSQL 16, which provides a number of new features and enhancements over version 15.

Notable enhancements include:

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

- You can now create custom configuration files and include them in the **pg_hba.conf** and **pg_ident.conf** files.
- PostgreSQL now supports regular expression matching on database and role entries in the **pg_hba.conf** file.

Other changes include:

- PostgreSQL is no longer distributed with the **postmaster** binary. Users who start the **postgresql** server by using the provided **systemd** unit file (the **systemctl start postgres** command) are not affected by this change. If you previously started the **postgresql** server directly through the **postmaster** binary, you must now use the **postgres** binary instead.
- PostgreSQL no longer provides documentation in PDF format within the package. Use the [online documentation](#) instead.

See also [Using PostgreSQL](#).

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

```
# yum module install postgresql:16
```

If you want to upgrade from an earlier **postgresql** stream within RHEL 8, follow the procedure described in [Switching to a later stream](#) and then migrate your PostgreSQL data as described in [Migrating to a RHEL 8 version of PostgreSQL](#).

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

[Jira:RHEL-3636](#)

Git rebased to version 2.43.0

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

Notable enhancements include:

- You can now use the new **--source** option with the **git check-attr** command to read the **.gitattributes** file from the provided tree-ish object instead of the current working directory.
- Git can now pass information from the **WWW-Authenticate** response-type header to credential helpers.
- In case of an empty commit, the **git format-patch** command now writes an output file containing a header of the commit instead of creating an empty file.
- You can now use the **git blame --contents=<file> <revision> -- <path>** command to find the origins of lines starting at **<file>** contents through the history that leads to **<revision>**.
- The **git log --format** command now accepts the **%(decorate)** placeholder for further customization to extend the capabilities provided by the **--decorate** option.

[Jira:RHEL-17103^{\[1\]}](#)

Git LFS rebased to version 3.4.1

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

Notable changes include:

- The **git lfs push** command can now read references and object IDs from standard input.
- Git LFS now handles alternative remotes without relying on Git.
- Git LFS now supports the **WWW-Authenticate** response-type header as a credential helper.

[Jira:RHEL-17102^{\[1\]}](#)

4.10. COMPILERS AND DEVELOPMENT TOOLS

elfutils rebased to version 0.190

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

- The **libelf** library now supports relative relocation (RELR).
- The **libdw** library now recognizes **.debug_[ct]u_index** sections.
- The **eu-readelf** utility now supports a new **-Ds, --use-dynamic --symbol** option to show symbols through the dynamic segment without using ELF sections.
- The **eu-readelf** utility can now show **.gdb_index** version 9.
- A new **eu-scrlines** utility compiles a list of source files associated with a specified DWARF or ELF file.
- A **debuginfod** server schema has changed for a 60% compression in file name representation (this requires reindexing).

[Jira:RHEL-15924](#)

valgrind updated to 3.22

The **valgrind** package has been updated to version 3.22. Notable improvements include:

- **valgrind memcheck** now checks that the values given to the C functions **memalign**, **posix_memalign**, and **aligned_alloc**, and the C++17 aligned **new** operator are valid alignment values.
- **valgrind memcheck** now supports mismatch detection for C++14 sized and C++17 aligned **new** and **delete** operators.
- Added support for lazy reading of DWARF debugging information, resulting in faster startup when **debuginfo** packages are installed.

[Jira:RHEL-15926](#)

Clang resource directory moved

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

[Jira:RHEL-9299](#)

A new grafana-selinux package

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

[Jira:RHEL-7503](#)

Updated GCC Toolset 13

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

Notable changes introduced in RHEL 8.10 include:

- The GCC compiler has been updated to version 13.2.1, which provides many bug fixes and enhancements that are available in upstream GCC.
- **binutils** now support AMD CPUs based on the **znver5** core through the **-march=znver5** compiler switch.
- **annobin** has been updated to version 12.32.
- The **annobin** plugin for GCC now defaults to using a more compressed format for the notes that it stores in object files, resulting in smaller object files and faster link times, especially in large, complex programs.

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

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

```
# 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, see [GCC Toolset 13](#).

[Jira:RHEL-25405^{\[1\]}](#)

LLVM Toolset rebased to version 17.0.6

LLVM Toolset has been updated to version 17.0.6.

Notable enhancements include:

- The opaque pointers migration is now completed.
- Removed support for the legacy pass manager in middle-end optimization.

Clang changes:

- C++20 coroutines are no longer considered experimental.
- Improved code generation for the **std::move** function and similar in unoptimized builds.

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

[Jira:RHEL-9028](#)

Rust Toolset rebased to version 1.75.0

Rust Toolset has been updated to version 1.75.0.

Notable enhancements include:

- Constant evaluation time is now unlimited
- Cleaner panic messages
- Cargo registry authentication
- **async fn** and opaque return types in traits

[Jira:RHEL-12964](#)

Go Toolset rebased to version 1.21.0

Go Toolset has been updated to version 1.21.0.

Notable enhancements include:

- **min**, **max**, and **clear** built-ins have been added.
- Official support for profile guided optimization has been added.
- Package initialization order is now more precisely defined.
- Type inferencing is improved.
- Backwards compatibility support is improved.

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

[Jira:RHEL-11872^{\[1\]}](#)

papi supports new processor microarchitectures

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

- AMD Zen 4
- 4th Generation Intel® Xeon® Scalable Processors

[Jira:RHEL-9336^{\[1\]}](#), [Jira:RHEL-9320](#), [Jira:RHEL-9337](#)

Ant rebased to version 1.10.9

The **ant:1.10** module stream has been updated to version 1.10.9. This version provides support for code signing, using a provider class and provider argument.



NOTE

The updated **ant:1.10** module stream provides only the **ant** and **ant-lib** packages. Remaining packages related to Ant are distributed in the **javapackages-tools** module in the unsupported CodeReady Linux Builder (CRB) repository and have not been updated.

Packages from the updated **ant:1.10** module stream cannot be used in parallel with packages from the **javapackages-tools** module. If you want to use the complete set of Ant-related packages, you must uninstall the **ant:1.10** module and disable it, [enable the CRB repository](#), and install the **javapackages-tools** module.

[Jira:RHEL-5365](#)

New package: maven-openjdk21

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

[Jira:RHEL-17126^{\[1\]}](#)

cmake rebased to version 3.26

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

- Added support for the C17 and C18 language standards.
- **cmake** can now query the **/etc/os-release** file for operating system identification information.
- Added support for the CUDA 20 and **nvtx3** libraries.
- Added support for the Python stable application binary interface.
- Added support for Perl 5 in the Simplified Wrapper and Interface Generator (SWIG) tool.

[Jira:RHEL-7396](#)

4.11. IDENTITY MANAGEMENT

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

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

If an IdP reference and an associated IdP user ID exist in IdM, you can use them to enable an IdM user to authenticate at the external IdP. After performing authentication and authorization at the external IdP, the IdM user receives a Kerberos ticket with single sign-on capabilities. The user must authenticate with the SSSD version available in RHEL 8.7 or later.

[Jira:RHELPLAN-123140^{\[1\]}](#)

ipa rebased to version 4.9.13

The **ipa** package has been updated from version 4.9.12 to 4.9.13. Notable changes include:

- The installation of an IdM replica now occurs against a chosen server, not only for Kerberos authentication but also for all IPA API and CA requests.
- The performance of the **cert-find** command has been improved dramatically for situations with a large number of certificates.
- The **ansible-freeipa** package has been rebased from version 1.11 to 1.12.1.

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

[Jira:RHEL-16936](#)

Deleting expired KCM Kerberos tickets

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

[Jira:SSSD-6216](#)

Support for bcrypt password hashing algorithm for local users

With this update, you can enable the **bcrypt** password hashing algorithm for local users. To switch to the **bcrypt** hashing algorithm:

1. Edit the **/etc/authselect/system-auth** and **/etc/authselect/password-auth** files by changing the **pam_unix.so sha512** setting to **pam_unix.so blowfish**.

2. Apply the changes:

```
█ # authselect apply-changes
```

3. Change the password for a user by using the **passwd** command.
4. In the **/etc/shadow** file, verify that the hashing algorithm is set to **\$2b\$**, indicating that the **bcrypt** password hashing algorithm is now used.

[Jira:SSSD-6790](#)

The idp Ansible module allows associating IdM users with external IdPs

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

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

[Jira:RHEL-16938](#)

IdM now supports the `idoverrideuser`, `idoverridegroup` and `idview` Ansible modules

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

idoverrideuser

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

idoverridegroup

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

idview

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

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

[Jira:RHEL-16933](#)

The delegation of DNS zone management enabled in **ansible-freeipa**

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

[Jira:RHEL-19133](#)

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

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

[Jira:RHEL-4963](#)

The `runasuser_group` parameter is now available in **ansible-freeipa ipasudorule**

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

[Jira:RHEL-19129](#)

389-ds-base rebased to version 1.4.3.39

The **389-ds-base** package has been updated to version 1.4.3.39.

[Jira:RHEL-19028](#)

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

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

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

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

[Jira:RHEL-19240](#)

samba rebased to version 4.19.4

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

- Command-line options in the **smbget** utility have been renamed and removed for a consistent user experience. However, this can break existing scripts or jobs that use the utility. See the **smbget --help** command and **smbget(1)** man page for further details about the new options.
- If the **winbind debug traceid** option is enabled, the **winbind** service now logs, additionally, the following fields:
 - **traceid**: Tracks the records belonging to the same request.
 - **depth**: Tracks the request nesting level.
- Samba no longer uses its own cryptography implementations and, instead, now fully uses cryptographic functionality provided by the GnuTLS library.
- The **directory name cache size** option was removed.

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

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

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

[Jira:RHEL-16483^{\[1\]}](#)

4.12. THE WEB CONSOLE

RHEL web console can now generate Ansible and shell scripts

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

[Jira:RHELDOS-17060^{\[1\]}](#)

Simplified managing storage and resizing partitions on Storage

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

[Jira:RHELDOCS-17056^{\[1\]}](#)

4.13. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The **ad_integration** RHEL system role now supports configuring dynamic DNS update options

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

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

You can change these and other settings using the new variables in **ad_integration**. For example, you can set **ad_dyndns_refresh_interval** to **172800** to change the DNS record refresh interval to 48 hours. For more details regarding the role variables, see the resources in the **/usr/share/doc/rhel-system-roles/ad_integration/** directory.

[Jira:RHELDOCS-17372^{\[1\]}](#)

The **metrics** RHEL System Role now supports configuring PMIE webhooks

With this update, you can automatically configure the **global_webhook_endpoint** PMIE variable using the **metrics_webhook_endpoint** variable for the **metrics** RHEL System Role. This enables you to provide a custom URL for your environment that receives messages about important performance events, and is typically used with external tools such as Event-Driven Ansible.

[Jira:RHEL-18170](#)

The **bootloader** RHEL system role

This update introduces the **bootloader** RHEL system role. You can use this feature for stable and consistent configuration of bootloaders and kernels on your RHEL systems. For more details regarding requirements, role variables, and example playbooks, see the README resources in the **/usr/share/doc/rhel-system-roles/bootloader/** directory.

[Jira:RHEL-3241](#)

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

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

[Jira:RHEL-15440](#)

Support for new **ha_cluster** System Role features

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

- Enablement of the repositories containing resilient storage packages, such as **dlm** or **gfs2**. A Resilient Storage subscription is needed to access the repository.
- Configuration of fencing levels, allowing a cluster to use multiple devices to fence nodes.

- Configuration of node attributes.

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

[Jira:RHEL-4624^{\[1\]}](#), [Jira:RHEL-22108](#), [Jira:RHEL-14090](#)

New RHEL System Role for configuring **fapolicyd**

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

[Jira:RHEL-16542](#)

The **network** RHEL System role now supports new route types

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

- **blackhole**
- **prohibit**
- **unreachable**

[Jira:RHEL-21491^{\[1\]}](#)

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

You can now configure or update the display name of the system registered to Red Hat Insights by using the new **rhc_insights.display_name** parameter. The parameter allows you to name the system based on your preference to easily manage systems in the Insights Inventory. If your system is already connected with Red Hat Insights, use the parameter to update the existing display name. If the display name is not set explicitly on registration, it is set to the hostname by default. It is not possible to automatically revert the display name to the hostname, but it can be set so manually.

[Jira:RHEL-16965](#)

The RHEL system roles now support LVM snapshot management

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

[Jira:RHEL-16553](#)

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

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

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

[Jira:RHEL-18963](#)

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

You can now configure or update the Ansible hostname for the systems registered to Red Hat Insights by using the new **rhc_insights.ansible_host** parameter. When set, the parameter changes the

ansible_host configuration in the `/etc/insights-client/insights-client.conf` file to your selected Ansible hostname. If your system is already connected with Red Hat Insights, this parameter will update the existing Ansible hostname.

[Jira:RHEL-16975](#)

ForwardToSyslog flag is now supported in the journald system role

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

[Jira:RHEL-21123](#)

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

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

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

[Jira:RHEL-19047](#)

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

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

[Jira:RHEL-15038](#)

The ad_integration RHEL System Role now supports custom SSSD settings

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

[Jira:RHEL-21134](#)

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

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

[Jira:RHEL-17667](#)

New logging_preserve_fqdn variable for the logging RHEL System Role

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

[Jira:RHEL-15933](#)

Support for creation of volumes without creating a file system

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

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

[Jira:RHEL-16213](#)

The **rhc** system role now supports RHEL 7 systems

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

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

[Jira:RHEL-16977](#)

New **mssql_ha_prep_for_pacemaker** variable

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

[Jira:RHEL-19204](#)

The **sshd** role now configures certificate-based SSH authentications

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

[Jira:RHEL-5985](#)

selinux role now supports configuring SELinux in disabled mode

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

[Jira:RHEL-15871](#)

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

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

[Jira:RHEL-19044](#)

4.14. VIRTUALIZATION

RHEL now supports Multi-FD migration of virtual machines

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

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

[Jira:RHELDPCS-16970^{\[1\]}](#)

Secure Execution VMs on IBM Z now support cryptographic coprocessors

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

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

[Jira:RHEL-11597^{\[1\]}](#)

You can now replace SPICE with VNC in the web console

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

Because the support for the SPICE protocol is deprecated in RHEL 8 and will be removed in RHEL 9, VMs that use the SPICE protocol fail to migrate to RHEL 9. However, RHEL 8 VMs use SPICE by default, so you must switch from SPICE to VNC for a successful migration.

[Jira:RHELDPCS-18289^{\[1\]}](#)

New virtualization features in the RHEL web console

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

- Add an SSH public key during virtual machine (VM) creation. This public key will be stored in the `~/.ssh/authorized_keys` file of the designated non-root user on the newly created VM, which provides you with an immediate SSH access to the specified user account.
- Select a **pre-formatted block device** type when creating a new storage pool. This is a more robust alternative to a **physical disk device** type, as it prevents unintentional reformatting of a raw disk device.

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

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

[Jira:RHELDPCS-18323^{\[1\]}](#)

4.15. RHEL IN CLOUD ENVIRONMENTS

New cloud-init clean option for deleting generated configuration files

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

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

Jira:RHEL-7312^[1]

RHEL instances on EC2 now support IPv6 IMDS connections

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

Jira:RHEL-7278

4.16. CONTAINERS

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. Notable bug fixes and enhancements over the previous version include:

Notable changes in Podman v4.9:

- You can now use Podman to load the modules on-demand by using the **podman --module <your_module_name>** command and to override the system and user configuration files.
- A new **podman farm** command with a set of the **create**, **set**, **remove**, and **update** subcommands has been added. With these commands, you can farm out builds to machines running podman for different architectures.
- A new **podman-compose** command has been added, which runs Compose workloads by using an external compose provider such as Docker compose.
- The **podman build** command now supports the **--layer-label** and **--cw** options.
- The **podman generate systemd** command is deprecated. Use Quadlet to run containers and pods under **systemd**.
- The **podman build** command now supports **Containerfiles** with the HereDoc syntax.
- The **podman machine init** and **podman machine set** commands now support a new **--usb** option. Use this option to allow USB passthrough for the QEMU provider.
- The **podman kube play** command now supports a new **--publish-all** option. Use this option to expose all containerPorts on the host.

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

Jira:RHELPLAN-167794^[1]

Podman now supports containers.conf modules

You can use Podman modules to load a predetermined set of configurations. Podman modules are **containers.conf** files in the Tom's Obvious Minimal Language (TOML) format.

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

- For rootless users: **\$HOME/.config/containers/containers.conf.modules**
- For root users: **/etc/containers/containers.conf.modules**, or **/usr/share/containers/containers.conf.modules**

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

- You can specify modules multiple times by using the **--module** option.
- If **<your_module_name>** is the absolute path, the configuration file will be loaded directly.
- The relative paths are resolved relative to the three module directories mentioned previously.
- Modules in **\$HOME** override those in the **/etc/** and **/usr/share/** directories.

For more information, see the [upstream documentation](#).

Jira:RHELPLAN-167830^[1]

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

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

Jira:RHELPLAN-167822^[1]

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

With Podman v4.9, the SQLite database backend for Podman, previously available as Technology Preview, is now fully supported. The SQLite database provides better stability, performance, and consistency when working with container metadata. The SQLite database backend is the default backend for new installations of RHEL 8.10. If you upgrade from a previous RHEL version, the default backend is BoltDB.

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

Jira:RHELPLAN-168179^[1]

Administrators can set up isolation for firewall rules by using nftables

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

Jira:RHELDPCS-16955^[1]

Containerfile now supports multi-line instructions

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

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

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

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

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

Jira:RHELPLAN-168184^[1]

Toolbx is now available

With Toolbx, you can install the development and debugging tools, editors, and Software Development Kits (SDKs) into the Toolbx fully mutable container without affecting the base operating system. The Toolbx container is based on the **registry.access.redhat.com/ubi8.10/toolbox:latest** image.

Jira:RHELDPCS-16241^[1]

CHAPTER 5. 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 5.1. 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	528482304
CONFIG_BPF	y
CONFIG_BPF_SYSCALL	y
CONFIG_HAVE_EBPF_JIT	y
CONFIG_BPF_JIT	y
CONFIG_BPF_JIT_ALWAYS_ON	y
CONFIG_DEBUG_INFO_BTF	y
CONFIG_DEBUG_INFO_BTF_MODULES	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 5.2. 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_sk_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_socket, bpf_skc_to_tcp_socket, bpf_skc_to_tcp_timewait_socket, bpf_skc_to_tcp_request_socket, bpf_skc_to_udp6_socket, 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_skb_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_wri table	bpf_map_lookup_elem, bpf_map_update_elem, bpf_map_delete_elem, bpf_probe_read, bpf_ktime_get_ns, bpf_get_prandom_u32, bpf_get_smp_processor_id, bpf_tail_call, bpf_get_current_pid_tgid, bpf_get_current_uid_gid, bpf_get_current_comm, bpf_perf_event_read, bpf_perf_event_output, bpf_get_stackid, bpf_get_current_task, bpf_current_task_under_cgroup, bpf_get_numa_node_id, bpf_probe_read_str, bpf_perf_event_read_value, bpf_get_stack, bpf_get_current_cgroup_id, bpf_map_push_elem, bpf_map_pop_elem, bpf_map_peek_elem, bpf_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_socket	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 5.3. 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 6. BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 8.10 that have a significant impact on users.

6.1. INSTALLER AND IMAGE CREATION

Installer now accepts additional time zone definitions in Kickstart files

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

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

Note: This change only applies to the **timezone** Kickstart command. The time zone selection in the graphical and text-based interactive interfaces remains unchanged. Existing Kickstart files for Red Hat Enterprise Linux 8 that had valid time zone selections do not require any updates.

[Jira:RHEL-13151^{\[1\]}](#)

6.2. SECURITY

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

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

[Jira:RHEL-9981^{\[1\]}](#)

SELinux policy allows **staff_r** confined users to run **sudo crontab**

Previously, the SELinux policy did not contain rules to allow confined users to run the **sudo crontab** command. As a consequence, confined users in the **staff_r** role could not use **sudo crontab** to edit other users' **crontab** schedules. This update adds a rule to the policy, and as a result, **staff_r** users can use **sudo crontab** to edit other users' **crontab** schedules.

[Jira:RHEL-1388](#)

SELinux policy contains rules for additional services and applications

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

- **sudo traceroute**
- **sudo tcpdump**
- **sudo dnf**

[Jira:RHEL-15398](#), [Jira:RHEL-1679](#), [Jira:RHEL-9947](#)

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

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

[Jira:RHEL-1628](#)

SELinux policy allows `rsyslogd` to execute confined commands

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

[Jira:RHEL-10087](#)

Large SSHD configuration files no longer prevent login

Previously, when the SSHD configuration file was larger than 256 KB, an error occurred when logging into the system. As a consequence, remote systems were unreachable. This update removes the file size limitation, and therefore users can log in to the system when the SSHD configuration file is larger than 256 KB.

[Jira:RHEL-5279](#)

6.3. SOFTWARE MANAGEMENT

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

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

[Jira:RHEL-17356](#)

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

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

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

[Jira:RHEL-10720](#)

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

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

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

[Jira:RHEL-6421](#)

6.4. SHELLS AND COMMAND-LINE TOOLS

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

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

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

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

[Jira:RHEL-24729^{\[1\]}](#)

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

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

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

The kernel log displayed either of the following messages:

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

```
alignment check failed: sunit/swidth vs. agsize
```

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

[Jira:RHEL-17354^{\[1\]}](#)

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

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

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

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

or

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

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

Jira:RHEL-17353^[1]

The **pkla-compact** binary is executed when the **polkit** is called on the **logind-session-monitor** event

Previously, re-verification of the authorizations for **polkit** actions was triggered by any **logind-session-monitor** event for all users. Each **CheckAuthorization** request executes the **polkit-pkla-compact** binary to check for legacy **.pkla** configuration files even if no such files are present on the system, which causes CPU usage to increase by the **polkit** daemon.

Currently, only the **logind-session** changes that are relevant for the **polkit** actions are observed. If the session's state changes, the **polkit** objects associated with the session trigger re-verification (**CheckAuthorization**). You must restart (**log out to login screen and re-login** or **reboot**) the **gnome-shell** for a successful update.

The **polkit-pkla-compact** binary is now a soft dependency. As a result, you can reduce the CPU intensity by uninstalling the **polkit-pkla-compact** binary only if there are no **.pkla** files present in **/etc/polkit-1/localauthority**, **/etc/polkit-1/localauthority.conf.d**, **/var/lib/polkit-1/localauthority** and their respective sub directories.

Jira:RHEL-34022^[1]

6.5. KERNEL

crash rebased to version 8.0.4

The **crash** utility has been upgraded to version 8.0.4, which provides multiple bug fixes. Notable fixes include:

- Fixed a segmentation fault when non-panicking CPUs failed to stop during a kernel panic.
- Fixed a critical error incorrectly preventing the kernel from panicking when the **panic_on_oops** kernel parameter was disabled.
- Fixed the **crash** utility resolving hashed freelist pointers for the kernel compiled with the **CONFIG_SLAB_FREELIST_HARDENED=y** configuration option.

- A change in the kernel module memory layout terminology replaced **module_layout** with **module_memory** to better indicate memory-related aspects of the **crash** utility. Prior to this change, the **crash** utility could not start a session and returned an error message like this:

```
crash: invalid structure member offset: module_core_size
FILE: kernel.c LINE: 3787 FUNCTION: module_init()
```

[Jira:RHEL-9010](#)

tuna launches GUI when needed

Previously, if you ran the **tuna** utility without any subcommand, it would launch the GUI. This behavior was desirable if you had a display. In the opposite case, **tuna** on a machine without a display would not exit gracefully. With this update, **tuna** detects whether you have a display, and the GUI is launched or not launched accordingly.

[Jira:RHEL-19179^{\[1\]}](#)

6.6. FILE SYSTEMS AND STORAGE

Multipathd now checks if a device is incorrectly queuing I/O

Previously, a multipath device restarted queuing I/O, even though it was configured to fail, under the following conditions:

- The multipath device was configured with the **queue_if_no_paths** parameter set to a number of retries.
- A path device was removed from the multipath device that had no working paths and was no longer queuing I/O.

With this update, the issue has been fixed. As a result, multipath devices no longer restarts queuing I/O if the queuing is disabled and a path is removed while there are no usable paths.

[Jira:RHEL-16563^{\[1\]}](#)

The **no_read_workqueue**, **no_write_workqueue**, and **try_verify_in_tasklet** options of the **dm-crypt** and **dm-verity** devices are temporarily disabled

Previously, the **dm-crypt** devices created by using either the **no_read_workqueue** or **no_write_workqueue** option and **dm-verity** devices created by using the **try_verify_in_tasklet** option caused memory corruption. Consequently, random kernel memory was corrupted, which caused various system problems. With this update, these options are temporarily disabled. Note that this fix can cause **dm-verity** and **dm-crypt** to perform slower on some workloads.

[Jira:RHEL-22232^{\[1\]}](#)

6.7. HIGH AVAILABILITY AND CLUSTERS

Issues with moving and banning clone and bundle resources now corrected

This bug fix addresses two limitations of moving bundled and clone resources:

- When a user tried to move a bundled resource out of its bundle or ban it from running in its bundle, **pcs** created a constraint but the constraint had no effect. This caused the move to fail

with an error message. With this fix, **pcs** disallows moving and banning bundled resources from their bundles and prints an error message noting that bundled resources cannot be moved out of their bundles.

- When a user tried to move a bundle or clone resource, **pcs** exited with an error message noting that bundle or clone resources cannot be moved. This fix relaxes validation of move commands. It is now possible to move clone and bundle resources. When moving clone resources, you must specify a destination node if more than one instance of a clone is running. Only one-replica bundles can be moved.

[Jira:RHEL-7584](#)

Output of **pcs status** command no longer shows warning for expired constraints

Previously, when moving a cluster resource created a temporary location constraint, the **pcs status** command displayed a warning even after the constraint expired. With this fix, the **pcs status** command filters out expired constraints and they no longer generate a warning message in the command output.

[Jira:RHEL-7668](#)

Disabling the **auto_tie_breaker** quorum option no longer allowed when SBD fencing requires it

Previously, **pcs** allowed a user to disable the **auto_tie_breaker** quorum option even when a cluster configuration required this option for SBD fencing to work correctly. With this fix, **pcs** generates an error message when a user attempts to disable **auto_tie_breaker** on a system where SBD fencing requires that the **auto_tie_breaker** option be enabled.

[Jira:RHEL-7731](#)

Configuring the **tls** and **keep_active_partition_tie_breaker** quorum device options without specifying **--force**

Previously, when configuring a quorum device, a user could not configure the **tls** and **keep_active_partition_tie_breaker** options for a quorum device model **net** without specifying the **--force** option. With this update, configuring these options no longer requires you to specify **--force**.

[Jira:RHEL-7745](#)

6.8. COMPILERS AND DEVELOPMENT TOOLS

ldconfig no longer crashes after an interrupted system upgrade

Previously, the **ldconfig** utility terminated unexpectedly with a segmentation fault when processing incomplete shared objects left in the **/usr/lib64** directory after an interrupted system upgrade. With this update, **ldconfig** ignores temporary files written during system upgrades. As a result, **ldconfig** no longer crashes after an interrupted system upgrade.

[Jira:RHEL-13720](#)

Improved **glibc** compatibility with applications using **dlclose** on shared objects involved in a dependency cycle

Previously, when unloading a shared object in a dependency cycle using the **dlclose** function in **glibc**, that object's ELF destructor might not have been called before all other objects were unloaded. As a consequence of this late ELF destructor execution, applications experienced crashes and other errors due to the initial shared object's dependencies already being deinitialized.

With this update, **glibc** has been fixed to first call the ELF destructor of the immediate object being unloaded before any other ELF destructors are executed. As a result, compatibility with applications using **dlclose** on shared objects involved in a dependency cycle is improved and crashes no longer occur.

[Jira:RHEL-1048^{\[1\]}](#)

Improved **glibc** wide-character write performance

Previously, the wide **stdio** stream implementation in **glibc** did not treat the default buffer size as large enough for wide-character write operations and used a 16-byte fallback buffer instead, negatively impacting performance. With this update, buffer management is fixed and the entire write buffer is used. As a result, **glibc** wide-character write performance is improved.

[Jira:RHEL-19824^{\[1\]}](#)

6.9. IDENTITY MANAGEMENT

Automembership plug-in no longer cleans up groups by default

Previously, the automember rebuild task first removed all the memberships values and then rebuilt the memberships from scratch. As a result, the rebuild task was expensive, especially if other **be_txn** plugins were enabled.

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

- Only one rebuilt task is allowed at a time.
- The Automembership plug-in no longer cleans up previous members by default. Use the new **--cleanup** CLI option to intentionally clean up memberships before rebuilding from scratch:

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

- Improved logging to display fixup progress.

[Jira:RHEL-5390^{\[1\]}](#)

Allocated memory now released when an operation is completed

Previously, memory allocated by the KCM for each operation was not being released until the connection was closed. As a result, for client applications that opened a connection and ran many operations on the same connection, it led to a noticeable memory increase because the allocated memory was not released until the connection closed. With this update, the memory allocated for an operation is now released as soon as the operation is completed.

[Jira:SSSD-7015](#)

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

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

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

[Jira:SSSD-6096](#)

SSSD correctly returns an error if no grace logins remain while changing a password

Previously, if a user's LDAP password had expired, SSSD tried to change the password even after the initial bind of the user failed as there were no more grace logins left. However, the error returned to the user did not indicate the reason for the failure. With this update, the request to change the password is aborted if the bind fails and SSSD returns an error message indicating there are no more grace logins and the password must be changed by another means.

[Jira:SSSD-6184](#)

Removing systems from a domain using the `realm leave` command

Previously, if multiple names were set for the `ad_server` option in the `sssd.conf` file, running the `realm leave` command resulted in parsing errors and the system was not removed from the domain. With this update, the `ad_server` option is properly evaluated and the correct domain controller name is used and the system is correctly removed from the domain.

[Jira:SSSD-6081](#)

KCM logs to the correct `sssd.kcm.log` file

Previously, `logrotate` correctly rotated the Kerberos Credential Manager (KCM) log files but KCM incorrectly wrote the logs to the old log file, `sssd_kcm.log.1`. If KCM was restarted, it used the correct log file. With this update, after `logrotate` is invoked, log files are rotated and KCM correctly logs to the `sssd_kcm.log` file.

[Jira:SSSD-6652](#)

The `realm leave --remove` command no longer asks for credentials

Previously, the `realm` utility did not correctly check if a valid Kerberos ticket was available when running the `realm leave` operation. As a result, users were asked to enter a password even though a valid Kerberos ticket was available. With this update, `realm` now correctly verifies if there is a valid Kerberos ticket and no longer requests the user to enter a password when running the `realm leave --remove` command.

[Jira:SSSD-6425](#)

IdM Vault encryption and decryption no longer fails in FIPS mode

Previously, IdM Vault used OpenSSL RSA-PKCS1v15 as the default padding wrapping algorithm. However, none of the FIPS certified modules in RHEL supported PKCS#1 v1.5 as a FIPS approved algorithm, causing IdM Vault to fail in FIPS mode. With this update, IdM Vault supports the RSA-OAEP padding wrapping algorithm as a fallback. As a result, IdM Vault encryption and decryption now work correctly in FIPS mode.

[Jira:RHEL-12153^{\[1\]}](#)

Non-CA IdM replica installation no longer fails with server affinity configured

In some scenarios, installing an IdM replica without a certificate authority (CA) failed with `CA_REJECTED` errors. The failure occurred due to the `certmonger` service attempting to retrieve certificates and resulted in incomplete replication details when adding a new replica to a complex

topology.

With this update, the IdM replica installation process happens against a specific IdM server that provides the necessary services such as Kerberos authentication and IdM API and CA requests. This ensures complete replication details when adding a new replica.

[Jira:RHEL-4964](#)

Kerberos Key Distribution Centers version 1.20 and later now process tickets generated from KDCs running version 1.18.2 and earlier

Previously, a compatibility issue occurred between a Key Distribution Center (KDC) running Kerberos version 1.20 or later and a KDC running version 1.18.2 or earlier. As a consequence, when evidence tickets issued by the KDC running Kerberos 1.20 or later were sent to the KDC running Kerberos 1.18.2 or earlier, the older KDC rejected the ticket granting service request because it lacked support for the **AD-SIGN****TICKET** attribute.

With this update, earlier versions of KDC now accept evidence tickets generated by KDCs running Kerberos 1.20 and newer, as they no longer require **AD-SIGN****TICKET** when a Privileged Attribute Certificate (PAC) is present.

[Jira:RHEL-10495](#)

SELinux labeling for `dirsrv` files was moved to `DEBUG` log level

Previously, SELinux labeling for `dirsrv` files had the **INFO** log level. With this update, the **DEBUG** log level is used for the `dirsrv` files the same way as it was in previous versions.

[Jira:RHEL-5143](#)

Directory Server no longer causes a segmentation fault when a backend is configured without a related suffix

Previously, if a backend was configured without a related suffix, Directory Server had a segmentation fault during startup. With this update, Directory Server checks if the suffix is associated with the backend before trying to access the suffix. As a result, the segmentation fault no longer occurs.

[Jira:RHEL-5107](#)

Directory Server no longer fails after abandoning the paged result search

Previously, a race condition was a reason for heap corruption and Directory Server failure during abandoning paged result search. With this update, the race condition was fixed, and Directory Server failure no longer occurs.

[Jira:RHEL-16338](#)

Directory Server now starts correctly after an upgrade if you configured a custom value for the connection table size

Previously, if you set a custom value for the connection table size and the `nsslapd-conntablesiz` attribute was present in the `dse.ldif` file, Directory Server did not start after an upgrade. With this release, Directory Server starts correctly after the upgrade with `nsslapd-conntablesiz` present in the `dse.ldif` file.

[Jira:RHEL-14025](#)

Directory Server no longer fails when Content Synchronization plug-in is enabled dynamically

Previously, if the Content Synchronization plug-in was enabled dynamically, the post-operation plug-in callback caused a segmentation fault because the pre-operation callback was not registered. With this update, the post-operation plug-in callback verifies that the memory is initialized and Directory Server no longer fails.

[Jira:RHEL-5135](#)

6.10. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Cluster start no longer times out when the SBD `delay-start` value is high

Previously, when a user configured SBD fencing in a cluster by using the `ha_cluster` System Role and set the `delay-start` option to a value close to or higher than 90 seconds, the cluster start timed out. This is because the default `systemd` start timeout is 90 seconds, which the system reached before the SBD start delay value. With this fix, the `ha_cluster` System Role overrides the `sbd.service` start timeout in `systemd` so that it is higher than the value of `delay-start`. This allows the system to start successfully even with high values of the `delay-start` option.

[Jira:RHEL-4684^{\[1\]}](#)

network role validates routing rules with `0.0.0.0/0` or `::/0`

Previously, when the `from:` or `to:` settings were set to the `0.0.0.0/0` or `::/0` addresses in the routing rule, the `network` RHEL System Role failed to configure the routing rule and rejected the settings as invalid. With this update, the `network` role allows `0.0.0.0/0` and `::/0` for `from:` and `to:` in routing rule validation. As a result, the role successfully configures the routing rules without raising the validation errors.

[Jira:RHEL-16501](#)

The `ha_cluster` system role now correctly configures a firewall on a `qnetd` host

Previously, when a user configured a `qnetd` host and set the `ha_cluster_manage_firewall` variable to `true` by using the `ha_cluster` system role, the role did not enable high-availability services in the firewall. With this fix, the `ha_cluster` system role now correctly configures a firewall on a `qnetd` host.

[Jira:RHEL-17874](#)

`keylime_server` role correctly reports registrar service status

Previously, when the `keylime_server` role playbook provided incorrect information, the role incorrectly reported the start as successful. With this update, the role now correctly reports a failure when incorrect information is provided, and the timeout when waiting for opened ports has been reduced from approximately 300 seconds to approximately 30 seconds.

[Jira:RHEL-21946](#)

The `postgresql` RHEL System Role now installs the correct version of PostgreSQL

Previously, if you tried to run the `postgresql` RHEL System Role with the `postgresql_version: "15"` variable defined on a RHEL managed node, PostgreSQL version 13 was installed instead of version 15. This bug has been fixed, and the `postgresql` role installs the version set in the variable.

[Jira:RHEL-21400](#)

The `podman` RHEL system role now sets and cancels linger properly for rootless containers

Previously, the **podman** RHEL System Role did not set and cancel `linger` properly for rootless containers. Consequently, deploying secrets or containers for rootless users produced errors in some cases, and failed to cancel `linger` when removing resources in some cases. With this update, the **podman** RHEL System Role ensures that `linger` is enabled for rootless users before doing any secret or container resource management, and ensures that `linger` is canceled for rootless users when there are no more secrets or container resources to be managed. As a result, the role correctly manages lingering for rootless users.

[Jira:RHEL-22228](#)

The **podman** RHEL system role now sets and cancels `linger` properly for rootless containers

Previously, the **podman** RHEL System Role did not set and cancel `linger` properly for rootless containers. Consequently, deploying secrets or containers for rootless users produced errors in some cases, and failed to cancel `linger` when removing resources in some cases. With this update, the **podman** RHEL System Role ensures that `linger` is enabled for rootless users before doing any secret or container resource management, and ensures that `linger` is canceled for rootless users when there are no more secrets or container resources to be managed. As a result, the role correctly manages lingering for rootless users.

[Jira:RHEL-22229](#)

Running read-scale clusters and installing **mssql-server-ha** no longer requires certain variables

Previously, if you used the **mssql** RHEL System Role to configure a read-scale cluster without certain variables (**`mssql_ha_virtual_ip`**, **`mssql_ha_login`**, **`mssql_ha_login_password`**, and **`mssql_ha_cluster_run_role`**), the role failed with an error message **Variable not defined**. However, these variables are not necessary to run a read-scale cluster. The role also tried to install the **mssql-server-ha**, which is not required for a read-scale cluster. With this fix, the requirement for these variables was removed. As a result, running a read-scale cluster proceeds successfully without the error message.

[Jira:RHEL-19202](#)

The **Kdump** system role works correctly when the `kexec_crash_size` file is busy

The `/sys/kernel/kexec_crash_size` file provides the size of the memory region allocated for crash kernel memory.

Previously, the **Kdump** system role failed when the `/sys/kernel/kexec_crash_size` file was busy. With this update, the system role retries reading the file when it is available. As a result, the system role no longer fails when the file is busy.

[Jira:RHEL-3354](#)

selinux role no longer uses the `item` loop variable

Previously, the **selinux** RHEL System Role used the `item` loop variable. This might have resulted in the following warning message when you called the **selinux** role from another role:

```
[WARNING]: TASK: fedora.linux_system_roles.selinux : Restore SELinux labels on filesystem tree:
The loop variable 'item' is already in use.
You should set the `loop_var` value in the `loop_control` option for the task to something else to avoid
variable collisions and unexpected behavior.
```

With this release, the **selinux** role uses `__selinux_item` as a loop variable. As a result, the warning that the **item** variable is already in use is no longer displayed even if you call the **selinux** role from another role.

[Jira:RHEL-19042](#)

Secret data is no longer logged with verbose logging

Previously, some tasks that handle secret data would log the contents. As a consequence, the logs showed secret data if verbose logging was being used. This update adds the **no_log: true** directive to tasks that can log secret data. As a result, secret data is not logged with verbose logging.

[Jira:RHEL-19242](#)

A volume quadlet service name no longer fails

Previously, starting the volume service name produced an error similar to the following one:

```
Could not find the requested service NAME.volume: host
```

With this update, the volume quadlet service name is changed to **basename-volume.service**. As a result, the volume service starts with no errors.

For more information, see [Volume unit](#) man page.

[Jira:RHEL-21402](#)

nbde_server role now works with socket overrides

Previously, the **nbde_server** RHEL System Role assumed that the only file in the **tangd** socket override directory was the **override.conf** file for a custom port. Consequently, the role deleted the directory if there was no port customization without checking other files, and the system re-created the directory in subsequent runs.

With this release, the role has been fixed to prevent changing attributes of the port override file and deleting the directory if there are other files. As a result, the role correctly works if **tangd** socket override files are managed also outside of the role.

[Jira:RHEL-25509](#)

6.11. VIRTUALIZATION

A dump failure no longer blocks IBM Z VMs with Secure Execution from running

Previously, when a dump of an IBM Z virtual machine (VM) with Secure Execution failed, the VM remained in a paused state and was blocked from running. For example, dumping a VM by using the **virsh dump** command fails if there is not enough space on the disk.

The underlying code has been fixed and Secure Execution VMs resume operation successfully after a dump failure.

[Jira:RHEL-16696^{\[1\]}](#)

CHAPTER 7. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 8.10.

For information on Red Hat scope of support for Technology Preview features, see [Technology Preview Features Support Scope](#).

7.1. 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](#)

7.2. 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 mirrored 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

7.3. 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 **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) subsystem in the Linux Kernel. Also, it configures devices through **sysfs** (pseudo-filesystem), saves and loads the configuration in the **json** format.

[Bugzilla:1843266^{\[1\]}](#)

7.4. 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]

7.5. 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^{\[1\]}](#)

7.6. 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](#)

7.7. 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:1667516, [Bugzilla:1724302](#), Bugzilla:1667225

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]

7.8. 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]

7.9. 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], Bugzilla:1501607, Jira:RHELPLAN-7677

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]

7.10. 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]

7.11. CONTAINERS

The **podman-machine** command is unsupported

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

Jira:RHELDPCS-16861^[1]

Building multi-architecture images is available as a Technology Preview

The **podman farm build** command, which you can use to create multi-architecture container images, is available as a Technology Preview.

A farm is a group of machines that have a unix podman socket running in them. The nodes in the farm can have different machines of different architectures. The **podman farm build** command is faster than the **podman build --arch --platform** command.

You can use **podman farm build** to perform the following actions:

- Build an image on all nodes in a farm.
- Bundle nodes up into a manifest list.
- Execute the **podman build** command on all the farm nodes.
- Push the images to the registry specified by using the **--tag** option.

- Locally create a manifest list.
- Push the manifest list to the registry.
The manifest list contains one image per native architecture type that is present in the farm.

Jira:RHELPLAN-154435^[1]

CHAPTER 8. 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](#).

8.1. 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]

8.2. 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^{\[1\]}](#)

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](#)

RHEL 8 and 9 OpenSSL certificate and signing containers are now deprecated

The OpenSSL portable certificate and signing containers available in the **ubi8/openssl** and **ubi9/openssl** repositories in the Red Hat Ecosystem Catalog are now deprecated due to low demand.

Jira:RHELDPCS-17974^[1]

8.3. SUBSCRIPTION MANAGEMENT

The deprecated `--token` option of `subscription-manager register` will stop working at the end of November 2024

The deprecated `--token=<TOKEN>` option of the **subscription-manager register** command will no longer be a supported authentication method from the end of November 2024. The default entitlement server, **subscription.rhsm.redhat.com**, will no longer be allowing token-based authentication. As a consequence, if you use **subscription-manager register --token=<TOKEN>**, the registration will fail with the following error message:

Token authentication not supported by the entitlement server

To register your system, use other supported authorization methods, such as including paired options `--username / --password` OR `--org / --activationkey` with the **subscription-manager register** command.

[Bugzilla:2170082](#)

8.4. 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](#)

8.5. SHELLS AND COMMAND-LINE TOOLS

Setting the `TMPDIR` variable in the ReaR configuration file is deprecated

Setting the **TMPDIR** environment variable in the `/etc/rear/local.conf` or `/etc/rear/site.conf` (ReaR configuration file), by using a statement such as **export TMPDIR=...**, is deprecated.

To specify a custom directory for ReaR temporary files, export the variable in the shell environment before executing ReaR. For example, execute the **export TMPDIR=...** statement and then execute the **rear** command in the same shell session or script.

Jira:RHELDPCS-18049^[1]

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

8.6. INFRASTRUCTURE SERVICES

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

8.7. NETWORKING

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 **cgdcboxd** package is deprecated

Control group data center bridging exchange daemon (**cgdcboxd**) 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 **cgdcboxd** package is deprecated and will be removed in the next major RHEL release.

[Bugzilla:2006665](#)

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.

[Bugzilla:2029338](#)

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.

[Bugzilla:2061288](#)

8.8. 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 a Technology Preview. Furthermore, 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 the 8.3 minor update of RHEL for Real Time 8, the diskless booting feature is no longer 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^{\[1\]}](#)

8.9. 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](#)

8.10. 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^{\[1\]}](#)

8.11. 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^{\[1\]}](#)

8.12. 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](#)

8.13. 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 **--allowerase** option to remove **gdb.i686** and install **gdb.x86_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](#)

8.14. 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]

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:RHELDPCS-17573^[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]

8.15. 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]

Several bitmap fonts have been deprecated

The following bitmap font packages have been deprecated:

- **bitmap-console-fonts**

- **bitmap-fixed-fonts**
- **bitmap-fonts-compatible**
- **bitmap-lucida-typewriter-fonts**

Bitmap fonts have a limited pixel size. When you try to set a font size that is unavailable, the text might display in a different size or a different font, possibly a scalable one. This also decreases the rendering quality of bitmap fonts and disrupts the user experience.

Additionally, the **fontconfig** system ignores the Portable Compiled Format (PCF), one of the major bitmap font formats, because it contains no metadata to estimate the language coverage.

Note that the **bitmap-fangsongti-fonts** bitmap font package continues to be supported as a dependency of the Lorax tool.

Jira:RHELDPCS-17623^[1]

8.16. 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]

8.17. 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]

8.18. RED HAT ENTERPRISE LINUX SYSTEM ROLES

The `network` System Role displays a deprecation warning when configuring teams on RHEL 9 nodes

The network teaming capabilities have been deprecated in RHEL 9. As a result, using the `network` RHEL System Role on 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](#)

The `mssql_ha_cluster_run_role` has been deprecated

The `mssql_ha_cluster_run_role` variable has been deprecated. Instead, use the `mssql_manage_ha_cluster` variable.

[Jira:RHEL-19203](#)

8.19. 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` or `virtio-vga` 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 SHA-2 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^{\[1\]}](#)

8.20. 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:RHELDPCS-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 **rhel8/openssl** has been deprecated

The **rhel8/openssl** container image has been deprecated.

Jira:RHELDPCS-18107^[1]

The **Inkscape** and **LibreOffice** Flatpak images are deprecated

The **rhel9/inkscape-flatpak** and **rhel9/libreoffice-flatpak** Flatpak images, which are available as Technology Previews, have been deprecated.

Red Hat recommends the following alternatives to these images:

- To replace **rhel9/inkscape-flatpak**, use the **inkscape** RPM package.
- To replace **rhel9/libreoffice-flatpak**, see the [LibreOffice deprecation release note](#) .

Jira:RHELDOCS-17102^[1]

pasta as a network name has been deprecated

The support for **pasta** as a network name value is deprecated and will not be accepted in the next major release of Podman, version 5.0. You can use the **pasta** network name value to create a unique network mode within Podman by employing the **podman run --network** and **podman create --network** commands.

Jira:RHELDOCS-17038^[1]

The BoltDB database backend has been deprecated

The BoltDB database backend is deprecated as of RHEL 8.10. In a future version of RHEL, the BoltDB database backend will be removed and will no longer be available to Podman. For Podman, use the SQLite database backend, which is now the default as of RHEL 8.10.

Jira:RHELDOCS-17461^[1]

The CNI network stack has been deprecated

The Container Network Interface (CNI) network stack is deprecated and will be removed in a future release. Use the Netavark network stack instead. For more information, see [Switching the network stack from CNI to Netavark](#).

Jira:RHELDOCS-17518^[1]

container-tools:4.0 has been deprecated

The **container-tools:4.0** module has been deprecated and will no longer receive security updates. To continue to build and run Linux Containers on RHEL, use the newer, stable, and supported module stream **container-tools:rhel8**.

For instructions on switching to a later stream, see [Switching to a later stream](#).

Jira:RHELPLAN-168223^[1]

8.21. 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
- arpwatch
- 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
- bash-doc
- 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
- bitmap-console-fonts
- bitmap-fixed-fonts
- bitmap-fonts-compatible
- bitmap-lucida-typewriter-fonts
- 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-condensed-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
- culmus-keteryg-fonts
- culmus-shofar-fonts
- custodia
- cyrus-imapd-vzic
- dbus-c++
- dbus-c++-devel
- dbus-c++-glib
- dbxtool
- dejavu-fonts-common
- 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
- 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
- dptfextract

- 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-osi
- eclipse-jdt
- eclipse-license
- eclipse-p2-discovery
- eclipse-pde
- eclipse-platform
- eclipse-swt
- ed25519-java
- ee4j-parent
- elfutils-devel-static
- elfutils-libelf-devel-static
- emacs-terminal
- emoji-picker
- 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
- gcc-toolset-12-annobin-annocheck
- gcc-toolset-12-annobin-docs
- gcc-toolset-12-annobin-plugin-gcc
- gcc-toolset-12-binutils

- gcc-toolset-12-binutils-devel
- gcc-toolset-12-binutils-gold
- 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-droid-kufi-fonts
- google-gson
- google-noto-kufi-arabic-fonts
- google-noto-naskh-arabic-fonts
- google-noto-naskh-arabic-ui-fonts
- google-noto-nastaliq-urdu-fonts
- google-noto-sans-balinese-fonts

- `google-noto-sans-bamum-fonts`
- `google-noto-sans-batak-fonts`
- `google-noto-sans-buginese-fonts`
- `google-noto-sans-buhid-fonts`
- `google-noto-sans-canadian-aboriginal-fonts`
- `google-noto-sans-cham-fonts`
- `google-noto-sans-cuneiform-fonts`
- `google-noto-sans-cypriot-fonts`
- `google-noto-sans-gothic-fonts`
- `google-noto-sans-gurmukhi-ui-fonts`
- `google-noto-sans-hanunoo-fonts`
- `google-noto-sans-inscriptional-pahlavi-fonts`
- `google-noto-sans-inscriptional-parthian-fonts`
- `google-noto-sans-javanese-fonts`
- `google-noto-sans-lepcha-fonts`
- `google-noto-sans-limbu-fonts`
- `google-noto-sans-linear-b-fonts`
- `google-noto-sans-lisu-fonts`
- `google-noto-sans-mandaic-fonts`
- `google-noto-sans-meetei-mayek-fonts`
- `google-noto-sans-mongolian-fonts`
- `google-noto-sans-myanmar-fonts`
- `google-noto-sans-myanmar-ui-fonts`
- `google-noto-sans-new-tai-lue-fonts`
- `google-noto-sans-ogham-fonts`
- `google-noto-sans-ol-chiki-fonts`
- `google-noto-sans-old-italic-fonts`
- `google-noto-sans-old-persian-fonts`
- `google-noto-sans-oriya-fonts`

- `google-noto-sans-oriya-ui-fonts`
- `google-noto-sans-phags-pa-fonts`
- `google-noto-sans-rejang-fonts`
- `google-noto-sans-runic-fonts`
- `google-noto-sans-samaritan-fonts`
- `google-noto-sans-saurashtra-fonts`
- `google-noto-sans-sundanese-fonts`
- `google-noto-sans-syloti-nagri-fonts`
- `google-noto-sans-syriac-eastern-fonts`
- `google-noto-sans-syriac-estrangela-fonts`
- `google-noto-sans-syriac-western-fonts`
- `google-noto-sans-tagalog-fonts`
- `google-noto-sans-tagbanwa-fonts`
- `google-noto-sans-tai-le-fonts`
- `google-noto-sans-tai-tham-fonts`
- `google-noto-sans-tai-viet-fonts`
- `google-noto-sans-tibetan-fonts`
- `google-noto-sans-tifinagh-fonts`
- `google-noto-sans-ui-fonts`
- `google-noto-sans-yi-fonts`
- `google-noto-serif-bengali-fonts`
- `google-noto-serif-devanagari-fonts`
- `google-noto-serif-gujarati-fonts`
- `google-noto-serif-kannada-fonts`
- `google-noto-serif-malayalam-fonts`
- `google-noto-serif-tamil-fonts`
- `google-noto-serif-telugu-fonts`
- `gphoto2`
- `graphviz-ruby`

- 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
- hspell

- 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
- khmeros-fonts-common
- ksc
- kurdit-unikurd-web-fonts
- kyotocabinet-libs
- langtable-data
- ldapjdk-javadoc
- lensfun
- lensfun-devel
- lftp-scripts
- libaec
- libaec-devel
- libappindicator-gtk3
- libappindicator-gtk3-devel
- libatomic-static
- libavc1394
- libblocksruntime
- libcacard
- libcacard-devel
- libcgroup
- libcgroup-pam
- 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
- libicu-doc
- libIDL
- libIDL-devel
- libidn
- libiec61883
- libindicator-gtk3
- libindicator-gtk3-devel
- libiscsi-devel
- libjose-devel
- libkcc
- libkcc-common
- libkcc-data
- libldb-devel
- liblogging
- libluksmeta-devel
- libmalaga
- libmcpp

- libmemcached
- libmemcached-libs
- libmetalink
- libmodulemd1
- libmongocrypt
- libmtp-devel
- libmusicbrainz5
- libmusicbrainz5-devel
- libnbd-devel
- libnice
- libnice-gstreamer1
- liboauth
- liboauth-devel
- libpfm-static
- libpng12
- libpsm2-compat
- 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-lb
- 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
- overpass-mono-fonts
- pakchois
- pandoc
- paps-libs
- paranamer
- paratype-pt-sans-caption-fonts
- 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-pytml
- 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`
- `qperf`
- `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`
- `resteasy-javadoc`
- `rasm-gtk`
- `rpm-plugin-priorreset`
- `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
- shim-ia32
- shrinkwrap
- sil-padauk-book-fonts
- sisu-inject
- sisu-mojos
- sisu-plexus
- skkdic
- SLOF
- smc-anjalioldlipi-fonts
- smc-dyuthi-fonts
- smc-fonts-common
- smc-kalyani-fonts
- smc-raghmalayalam-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
- systemd-tests
- tcl-brlapi
- testng
- thai-scalable-laksaman-fonts
- tibetan-machine-uni-fonts
- timedatex
- torque-libs
- 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
- voikko-tools
- vorbis-tools
- weld-parent
- wodim
- woodstox-core
- wqy-microhei-fonts
- wqy-unibit-fonts
- xdelta
- xmlgraphics-commons
- xmlstreambuffer
- xinetd
- xorg-x11-apps
- xorg-x11-driv-qxl
- xorg-x11-server-Xspice
- xpp3

- xsane-gimp
- xsom
- xz-java
- xz-java-javadoc
- yajl-devel
- yp-tools
- ypbind
- ypserv
- zsh-html

8.22. 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 8.1. Deprecated devices

Device ID	Driver	Device name
	hns_roce	
	ebtables	
	arp_tables	
	ip_tables	
	ip6_tables	

Device ID	Driver	Device name
	ip6_set	
	ip_set	
	nft_com pat	
	usnic_ve rbs	
	vmw_pvr dma	
	hfi1	
	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	megarai d_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

Device ID	Driver	Device name
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
0x1000:0x0087	mpt3sas	Broadcom / LSI: SAS2308 PCI-Express Fusion-MPT SAS-2
	myri10g e	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_ r xe)	
	HNS- RoCE	HNS GE/10GE/25GE/50GE/100GE RDMA Network Controller
	liquidio	Cavium LiquidIO Intelligent Server Adapter Driver

Device ID	Driver	Device name
	liquidio_vf	Cavium LiquidIO Intelligent Server Adapter Virtual Function Driver

Table 8.2. Unmaintained devices

Device ID	Driver	Device name
	dl2k	
	dlci	
	dnet	
	hdlc_fr	
	rdma_rx e	
	nicvf	
	nicpf	
	siw	
	e1000	Intel® PRO/1000 Network Driver
	mptbase	Fusion MPT SAS Host driver
	mptsas	Fusion MPT SAS Host driver
	mptscsi h	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]

Device ID	Driver	Device name
	nvmet_t cp	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 9. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 8.10.

9.1. 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 work around 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](#)

9.2. SECURITY

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](#)

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, OSCP 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](#)

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](#)

Kickstart uses `org_fedora_oscaped` instead of `com_redhat_oscaped` in RHEL 8

The Kickstart references the Open Security Content Automation Protocol (OSCAP) Anaconda add-on as `org_fedora_oscaped` instead of `com_redhat_oscaped`, 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^{\[1\]}](#)

The `semanage fcontext` command reorders local modifications

The `semanage fcontext -l -C` command lists local file context modifications stored in the `file_contexts.local` file. The `restorecon` utility processes the entries in the `file_contexts.local` from the most recent entry to the oldest. However, `semanage fcontext -l -C` lists the entries in a different order.

This mismatch between processing order and listing order might cause problems when managing SELinux rules.

[Jira:RHEL-24461^{\[1\]}](#)

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](#)

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](#)

9.3. RHEL FOR EDGE

composer-cli fails to build RHEL for Edge images when nodejs or npm is included

Currently, while using RHEL image builder, you cannot customize your RHEL 8 Edge images with the **nodejs** and **npm** packages, because it is not possible to build a RHEL for Edge image with the **nodejs** package. The NPM package manager expects its configuration in the **{prefix}/etc/npmrc** directory and the npm RPM packages a symlink at the **/usr/etc/npmrc** directory pointing to **/etc/npmrc**. To work around this problem, install the **nodejs** and **npm** packages after building your RHEL for Edge system.

[Jira:RHELDPCS-17126^{\[1\]}](#)

9.4. 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](#)

9.5. SOFTWARE MANAGEMENT

YUM functionalities or plug-ins might log messages even if a logging service is not available

Certain **YUM** functionalities or plug-ins might log messages to standard output or standard error when a logging service is not available. The level of the log message indicates where the message is logged:

- Information messages are logged to standard output.
- Error and debugging messages are logged to standard error.

As a consequence, when scripting **YUM** options, unwanted log messages on standard output or standard error can affect the functionality of the script.

To work around this issue, suppress the log messages from standard output and standard error by using the **yum -q** command. This suppresses log messages but not command results that are expected on standard output.

Jira:RHELPLAN-50409^[1]

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

9.6. 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](#)

The ReaR rescue image on UEFI systems with Secure Boot enabled fails to boot with the default settings

ReaR image creation by using the **rear mkrescue** or **rear mkbackup** command fails with the following message:

```
grub2-mkstandalone may fail to make a bootable EFI image of GRUB2 (no /usr/*/grub*/x86_64-efi/moddep.lst file)
(...)
grub2-mkstandalone: error: /usr/lib/grub/x86_64-efi/modinfo.sh doesn't exist. Please specify --target or --directory.
```

The missing files are part of the **grub2-efi-x64-modules** package. If you install this package, the rescue image is created successfully without any errors. When the **UEFI** Secure Boot is enabled, the rescue image is not bootable because it uses a boot loader that is not signed.

To work around this problem, add the following variables to the **/etc/rear/local.conf** or **/etc/rear/site.conf** ReaR configuration file):

```
UEFI_BOOTLOADER=/boot/efi/EFI/redhat/grubx64.efi
SECURE_BOOT_BOOTLOADER=/boot/efi/EFI/redhat/shimx64.efi
```

With the suggested workaround, the image can be produced successfully even on systems without the **grub2-efi-x64-modules** package, and it is bootable on systems with Secure Boot enabled. In addition, during the system recovery, the bootloader of the recovered system is set to the **EFI** shim bootloader.

For more information about **UEFI, Secure Boot**, and **shim bootloader**, see the [UEFI: what happens when booting the system](#) Knowledge Base article.

Jira:RHELDOCS-18064^[1]

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

The %util and svctm columns produced by sar and iostat utilities are invalid

When you collect system usage statistics by using the **sar** or **iostat** utilities on a system with kernel version **4.18.0-55.el8** or later, the **%util** and **svctm** columns produced by **sar** or **iostat** might contain invalid data.

Jira:RHEL-23074^[1]

9.7. 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 brlty package is not multilib compatible

It is not possible to have both 32-bit and 64-bit versions of the **brlty** package installed. You can either install the 32-bit (**brlty.i686**) or the 64-bit (**brlty.x86_64**) version of the package. The 64-bit version is recommended.

Bugzilla:2008197

9.8. 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]

9.9. 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.

- 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: iwlfwif 0000:09:00.0: Failed to start RT ucode: -110
kernel: iwlfwif 0000:09:00.0: WRT: Collecting data: ini trigger 13 fired (delay=0ms)
kernel: iwlfwif 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]

9.10. 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:1798631, Bugzilla:1808012

System panics after enabling the IOMMU

Enabling the Input-Output Memory Management Unit (IOMMU) on the kernel command line by setting the **intel_iommu** parameter to **on** results in system panic with general protection fault for the **0x6b6b6b6b6b6b6b6b: 0000** non-canonical address.

To work around this problem, ensure that **intel_iommu** is set to **off**.

Jira:RHEL-1765^[1]

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

9.11. 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](#)

9.12. 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/sss/sss.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

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](#)

9.13. 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](#)

9.14. 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^{\[1\]}](#)

9.15. RED HAT ENTERPRISE LINUX SYSTEM ROLES

Using the RHEL system role with Ansible 2.9 can display a warning about using **dnf** with the **command** module

Since RHEL 8.8, the RHEL system roles no longer use the **warn** parameter in with the **dnf** module because this parameter was removed in Ansible Core 2.14. However, if you use the latest **rhel-system-roles** package still with Ansible 2.9 and a role installs a package, one of the following warnings can be displayed:

[WARNING]: Consider using the dnf module rather than running 'dnf'. If you need to use command because dnf is insufficient you can add 'warn: false' to this command task or set 'command_warnings=False' in ansible.cfg to get rid of this message.

[WARNING]: Consider using the yum, dnf or zypper module rather than running 'rpm'. If you need to use command because yum, dnf or zypper is insufficient you can add 'warn: false' to this command task or set 'command_warnings=False' in ansible.cfg to get rid of this message.

If you want to hide these warnings, add the **command_warnings = False** setting to the **[Defaults]** section of the **ansible.cfg** file. However, note that this setting disables all warnings in Ansible.

[Jira:RHELDPCS-17954](#)

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 work around 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](#)

Configuring the `imuxsock` input basics type causes a problem

Configuring the "imuxsock" input basics type through the `logging` RHEL system role and the `use_imuxsock` option cause a problem in the resulting configuration on the managed nodes. This role sets the `name` parameter, however, the "imuxsock" input type does not support the `name` parameter. As a result, the `rsyslog` logging utility prints the `parameter 'name' not known – typo in config file?` error.

[Jira:RHELDOCS-18326](#)

9.16. 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 **libvirt** 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](#)

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]

9.17. 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=xfst --grow --asprimary --size=8192 --ondisk=/dev/disk/by-id/wwn-
0x600224807270e09717645b1890f8a9a2
part /home --fstype="xfst" --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 -l 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^{\[1\]}](#)

9.18. 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.

[Jira:RHEL-2345](#)

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]

9.19. CONTAINERS

Running systemd within an older container image does not work

Running systemd within an older container image, for example, **centos:7**, does not work:

```
$ podman run --rm -ti centos:7 /usr/lib/systemd/systemd
Storing signatures
Failed to mount cgroup at /sys/fs/cgroup/systemd: Operation not permitted
[!!!!!!] Failed to mount API filesystems, freezing.
```

To work around this problem, use the following commands:

```
# mkdir /sys/fs/cgroup/systemd
# mount none -t cgroup -o none,name=systemd /sys/fs/cgroup/systemd
# podman run --runtime /usr/bin/crun --annotation=run.oci.systemd.force_cgroup_v1=/sys/fs/cgroup -
-rm -ti centos:7 /usr/lib/systemd/systemd
```

Jira:RHELPLAN-96940^[1]

CHAPTER 10. INTERNATIONALIZATION

10.1. 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

10.2. 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	Jira:RHEL-19028 , Jira:RHEL-19240 , Jira:RHEL-5390 , Jira:RHEL-5143 , Jira:RHEL-5107 , Jira:RHEL-16338 , Jira:RHEL-14025 , Jira:RHEL-5135
Release Notes	Jira:RHELDOCS-17954 , Jira:RHELDOCS-18326 , Jira:RHELDOCS-16861 , Jira:RHELDOCS-16755 , Jira:RHELDOCS-16612 , Jira:RHELDOCS-17102 , Jira:RHELDOCS-17518
SLOF	Bugzilla:1910848
accel-config	Bugzilla:1843266
anaconda	Jira:RHEL-13151 , Bugzilla:2050140 , Jira:RHEL-4707 , Jira:RHEL-4711 , Jira:RHEL-4744
ansible-collection-microsoft-sql	Jira:RHEL-19204 , Jira:RHEL-19202 , Jira:RHEL-19203
ansible-freeipa	Jira:RHEL-16938 , Jira:RHEL-16933 , Jira:RHEL-19133 , Jira:RHEL-4963 , Jira:RHEL-19129
ant	Jira:RHEL-5365
apr	Bugzilla:1819607
audit	Jira:RHEL-15001
authselect	Bugzilla:1892761
bacula	Jira:RHEL-6859
brlTTY	Bugzilla:2008197
chrony	Jira:RHEL-21069
clang	Jira:RHEL-9299
cloud-init	Jira:RHEL-7312 , Jira:RHEL-7278 , Jira:RHEL-12122
cmake	Jira:RHEL-7396
cockpit	Bugzilla:1666722

Component	Tickets
coreutils	Bugzilla:2030661
corosync-qdevice	Bugzilla:1784200
crash	Jira:RHEL-9010 , Bugzilla:1906482
crash-ptdump-command	Bugzilla:1838927
createrepo_c	Bugzilla:1973588
crypto-policies	Jira:RHEL-2345 , Bugzilla:1919155 , Bugzilla:1660839
device-mapper-multipath	Jira:RHEL-6677 , Jira:RHEL-16563 , Bugzilla:2022359 , Bugzilla:2011699
distribution	Jira:RHEL-17090 , Bugzilla:1657927
dnf	Bugzilla:1986657
dnf-plugins-core	Jira:RHEL-17356 , Jira:RHELPLAN-50409
edk2	Bugzilla:1741615 , Bugzilla:1935497
elfutils	Jira:RHEL-15924
fapolicyd	Jira:RHEL-520 , Bugzilla:2054741
fence-agents	Bugzilla:1775847
firewalld	Bugzilla:1871860
gcc-toolset-13-binutils	Jira:RHEL-25405
gdb	Bugzilla:1853140
git	Jira:RHEL-17103
git-lfs	Jira:RHEL-17102
glibc	Jira:RHEL-13720 , Jira:RHEL-10481 , Jira:RHEL-19824
gnome-shell-extensions	Bugzilla:1717947
gnome-software	Bugzilla:1668760
gnutls	Bugzilla:1628553

Component	Tickets
golang	Jira:RHEL-11872
grafana	Jira:RHEL-7503
grub2	Jira:RHEL-15856 , Jira:RHEL-15583 , Bugzilla:1583445
initscripts	Bugzilla:1875485
ipa	Jira:RHEL-16936 , Jira:RHEL-12153 , Jira:RHEL-4964 , Jira:RHEL-10495 , Jira:RHEL-4847 , Jira:RHEL-4898 , Bugzilla:1664719 , Bugzilla:1664718
ipmitool	Jira:RHEL-6846
kernel	Bugzilla:2041881 , Jira:RHEL-11597 , 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:1605216 , Bugzilla:1519039 , Bugzilla:1627455 , Bugzilla:1501618 , Bugzilla:1633143 , Bugzilla:1814836 , Bugzilla:1839311 , Bugzilla:1696451 , Bugzilla:1348508 , Bugzilla:1836977 , Bugzilla:1878207 , Bugzilla:1665295 , Bugzilla:1871863 , Bugzilla:1569610 , Bugzilla:1794513
kernel / DMA Engine	Jira:RHEL-10097
kernel / Networking / NIC Drivers	Jira:RHEL-11398
kernel / Networking / Protocol / tcp	Jira:RHEL-6113
kernel / Storage / Device Mapper / Crypt	Jira:RHEL-22232
kernel / Storage / Storage Drivers	Jira:RHEL-1765
kernel / Virtualization / KVM	Jira:RHEL-2451
kernel-rt / Other	Jira:RHEL-9318
kexec-tools	Bugzilla:2111855
kmod	Bugzilla:2103605
krb5	Jira:RHEL-4910 , Bugzilla:2125318 , Bugzilla:1877991

Component	Tickets
libdnf	Jira:RHEL-6421
libgnome-keyring	Bugzilla:1607766
libguestfs	Bugzilla:1554735
libkcapi	Jira:RHEL-5366 , Jira:RHEL-15300
librdkafka	Jira:RHEL-12892
librepo	Jira:RHEL-10720
libreswan	Bugzilla:1989050
libselinux-python-2.8-module	Bugzilla:1666328
libvirt	Bugzilla:1664592 , Bugzilla:1332758 , Bugzilla:2143160 , Bugzilla:1528684
linuxptp	Jira:RHEL-21326
llvm-toolset	Jira:RHEL-9028
lvm2	Bugzilla:1496229 , Bugzilla:1768536
mariadb	Jira:RHEL-3637 , Bugzilla:1942330
maven	Jira:RHEL-17126
mesa	Bugzilla:1886147
nfs-utils	Bugzilla:2081114 , Bugzilla:1592011
nginx	Jira:RHEL-14714
nispor	Bugzilla:2153166
nss	Bugzilla:1817533 , Bugzilla:1645153
nss_nis	Bugzilla:1803161
opencryptoki	Jira:RHEL-11413
opencv	Bugzilla:1886310

Component	Tickets
openmpi	Bugzilla:1866402
opencs	Jira:RHEL-4077 , Bugzilla:1947025
openscap	Bugzilla:2161499
openssh	Jira:RHEL-1684 , Jira:RHEL-5279 , Bugzilla:2044354
openssl	Jira:RHEL-17689 , Bugzilla:1810911
osbuild-composer	Jira:RHEL-4649
oscap-anaconda-addon	Jira:RHEL-1826 , Bugzilla:1843932 , Bugzilla:1834716 , Bugzilla:1665082 , Jira:RHEL-1810
papi	Jira:RHEL-9336
pcs	Jira:RHEL-7584 , Jira:RHEL-7668 , Jira:RHEL-7731 , Jira:RHEL-7745 , Bugzilla:1619620 , Bugzilla:1851335
perl-DateTime-TimeZone	Jira:RHEL-35685
php	Jira:RHEL-14705
pki-core	Bugzilla:1729215 , Jira:RHEL-13125 , Bugzilla:1628987
podman	Jira:RHELPLAN-167794 , Jira:RHELPLAN-167830 , Jira:RHELPLAN-167822 , Jira:RHELPLAN-168179 , Jira:RHELPLAN-168184 , Jira:RHELPLAN-154435 , Jira:RHELPLAN-168223
policycoreutils	Jira:RHEL-24461
polkit	Jira:RHEL-34022
postfix	Bugzilla:1711885
postgresql	Jira:RHEL-3636
pykickstart	Bugzilla:1637872
python3.11-lxml	Bugzilla:2157673
python36-3.6-module	Bugzilla:2165702

Component	Tickets
qemu-kvm	Jira:RHEL-11597 , Jira:RHEL-16696 , Jira:RHEL-13336 , Bugzilla:1740002 , Bugzilla:1719687 , Bugzilla:1966475 , Bugzilla:1792683 , Bugzilla:1651994
rear	Jira:RHEL-24729 , Jira:RHEL-17354 , Jira:RHEL-17353 , Bugzilla:1925531 , Bugzilla:2083301
redhat-support-tool	Bugzilla:2064575
restore	Bugzilla:1997366
rhel-system-roles	Jira:RHEL-18170 , Jira:RHEL-3241 , Jira:RHEL-15440 , Jira:RHEL-4624 , Jira:RHEL-16542 , Jira:RHEL-21491 , Jira:RHEL-16965 , Jira:RHEL-16553 , Jira:RHEL-18963 , Jira:RHEL-16975 , Jira:RHEL-21123 , Jira:RHEL-19047 , Jira:RHEL-15038 , Jira:RHEL-21134 , Jira:RHEL-14022 , Jira:RHEL-17667 , Jira:RHEL-15933 , Jira:RHEL-16213 , Jira:RHEL-16977 , Jira:RHEL-5985 , Jira:RHEL-4684 , Jira:RHEL-16501 , Jira:RHEL-17874 , Jira:RHEL-21946 , Jira:RHEL-21400 , Jira:RHEL-15871 , Jira:RHEL-22228 , Jira:RHEL-22229 , Jira:RHEL-3354 , Jira:RHEL-19042 , Jira:RHEL-19044 , Jira:RHEL-19242 , Jira:RHEL-21402 , Jira:RHEL-25509 , Bugzilla:2186908 , Bugzilla:2021685 , Bugzilla:2006081
rpm	Bugzilla:1688849
rsyslog	Bugzilla:1679512 , Jira:RHELPLAN-10431
rteval	Jira:RHEL-8967 , Jira:RHEL-21926
rtla	Jira:RHEL-10081
rust-toolset	Jira:RHEL-12964
samba	Jira:RHEL-16483 , Bugzilla:2009213 , Jira:RHELPLAN-13195
scap-security-guide	Jira:RHEL-25250 , Bugzilla:2028428 , Bugzilla:2118758 , Jira:RHEL-1804 , Jira:RHEL-1897
selinux-policy	Jira:RHEL-9981 , Jira:RHEL-1388 , Jira:RHEL-15398 , Jira:RHEL-1628 , Jira:RHEL-10087 , Bugzilla:2166153 , Bugzilla:1461914
sos	Bugzilla:2011413
spice	Bugzilla:1849563
sssd	Jira:SSSD-7015 , Bugzilla:2065692 , Bugzilla:2056483 , Bugzilla:1947671
sssd_kcm	Jira:SSSD-7015

Component	Tickets
stunnel	Jira:RHEL-2340
subscription-manager	Bugzilla:2170082
sysstat	Jira:RHEL-12008 , Jira:RHEL-23074
tuna	Jira:RHEL-19179
tuned	Bugzilla:2113900
udica	Bugzilla:1763210
valgrind	Jira:RHEL-15926
vdo	Bugzilla:1949163
virt-manager	Bugzilla:2026985
wayland	Bugzilla:1673073
webkit2gtk3	Jira:RHEL-4158
xorg-x11-server	Bugzilla:1698565

Component	Tickets
other	Jira:RHELDOCS-17369, Jira:RHELDOCS-17372, Jira:RHELDOCS-16955, Jira:RHELDOCS-16241, Jira:RHELDOCS-16970, Jira:RHELDOCS-17060, Jira:RHELDOCS-17056, Jira:RHELDOCS-16337, Jira:RHELDOCS-17261, Jira:RHELPLAN-123140, Jira:RHELDOCS-18289, Jira:RHELDOCS-18323, Jira:SSSD-6184, Bugzilla:2025814, Bugzilla:2077770, Bugzilla:1777138, Bugzilla:1640697, Bugzilla:1697896, Bugzilla:1961722, Jira:RHELDOCS-18064, Jira:RHELDOCS-18049, 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, Jira:RHELDOCS-17126, 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, Bugzilla:1932222, Bugzilla:1686057, Bugzilla:1748980, Jira:RHELPLAN-71200, Jira:RHELPLAN-45858, Bugzilla:1871025, Bugzilla:1871953, Bugzilla:1874892, Bugzilla:1916296, Jira:RHELDOCS-17573, 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:RHELDOCS-18107, Jira:RHELPLAN-153267, Bugzilla:2225332, Jira:RHELPLAN-147538, Jira:RHELDOCS-16612, Jira:RHELDOCS-17102, Jira:RHELDOCS-16300, Jira:RHELDOCS-17038, Jira:RHELDOCS-17461, Jira:RHELDOCS-17518, Jira:RHELDOCS-17623

APPENDIX B. REVISION HISTORY

0.0-2

Fri June 7 2024, Brian Angelica (bangelic@redhat.com)

- Updated a Known Issue in [Jira:RHELDOCS-18326](#) (Red Hat Enterprise Linux System Roles).

0.0-1

Thu May 23 2024, Brian Angelica (bangelic@redhat.com)

- Release of the Red Hat Enterprise Linux 8.10 Release Notes.

0.0-0

Wed March 27 2024, Lucie Vařáková (lvarakova@redhat.com)

- Release of the Red Hat Enterprise Linux 8.10 Beta Release Notes.