Red Hat Enterprise Linux 7
7.2 Release Notes

Release Notes for Red Hat Enterprise Linux 7.2

Red Hat Customer Content Services

Red Hat Enterprise Linux 7 7.2 Release Notes

Release Notes for Red Hat Enterprise Linux 7.2

Red Hat Customer Content Services

Legal Notice

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 7.2 and document known problems in this release. For detailed documentation on all changes to Red Hat Enterprise Linux for the 7.2 update, refer to Errata on the Red Hat Customer Portal.

Table of Contents

Preface

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security, and bug fix errata. The *Red Hat Enterprise Linux 7.2 Release Notes* document describes the major changes made to the Red Hat Enterprise Linux 7 operating system and its accompanying applications for this minor release, as well as known problems and a complete list of all currently available Technology Previews.

Capabilities and limits of Red Hat Enterprise Linux 7 as compared to other versions of the system are available in the Knowledge Base article available at https://access.redhat.com/articles/rhel-limits.

For information regarding the Red Hat Enterprise Linux life cycle, refer to https://access.redhat.com/support/policy/updates/errata/.

Chapter 1. Architectures

Red Hat Enterprise Linux 7.2 is available as a single kit on the following architectures: [1]

- 64-bit AMD
- 64-bit Intel
- IBM POWER7+ and POWER8 (big endian) [2]
- IBM POWER8 (little endian) [3]
- IBM System z [4]

---

[1] Note that the Red Hat Enterprise Linux 7.2 installation is supported only on 64-bit hardware. Red Hat Enterprise Linux 7.2 is able to run 32-bit operating systems, including previous versions of Red Hat Enterprise Linux, as virtual machines.

[2] Red Hat Enterprise Linux 7.2 (big endian) is currently supported only on PowerVM.

[3] Red Hat Enterprise Linux 7.2 (little endian) is currently supported on PowerVM and PowerHV (bare metal).

[4] Note that Red Hat Enterprise Linux 7.2 supports IBM zEnterprise 196 hardware or later; IBM System z10 mainframe systems are no longer supported and will not boot Red Hat Enterprise Linux 7.2.

Part I. New Features

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 7.2.

Chapter 2. General Updates

Cross channel package dependency improvements

Yum has been enhanced to prompt the end user to search disabled package repositories on the system when a package dependency error occurs. This change will allow users to quickly resolve dependency errors by first checking all known channels for the missing package dependency.

To enable this functionality, execute **`yum update yum subscription-manager`** prior to upgrading your machine to Red Hat Enterprise Linux 7.2.

Please see the System and Subscription Management chapter for further details on the implementation of this feature.

RELRO protection now properly applied when requested

Previously, binary files started by the system loader would, in some cases, lack the Relocation Read-Only (RELRO) protection even though this had been explicitly requested when the application was built. This was due to a miscommunication between the static linker and the system loader. The underlying source code of the linker has been adjusted to ensure that it makes it possible for the loader to apply the RELRO protection, thus restoring the security feature for applications. Applications and all dependent object files, archives, and libraries built with an alpha or beta version of *binutils* should be rebuilt to correct this defect. This update fixes the issue on the AMD64, Intel 64, 64-bit PowerPC, and 64-bit ARM architectures.

More diagnostic information and a renamed plug-in for sosreport

The sosreport tool has been enhanced to collect process-related information from various applications, including ptp, lastlog, and ethtool. As a part of this change, the **`startup`** plug-in has been renamed to **`services`** in order to better communicate its function.

Enable virtio network device renaming

This update adds a new persistent naming scheme for the virtio driver, which enables virtio network device renaming. To enable this feature in Red Hat Enterprise Linux 7.2, add the **`net.ifnames=1`** kernel parameter while booting.

Support for DIF/DIX (T10 PI) on specified hardware

SCSI T10 DIF/DIX is fully supported in Red Hat Enterprise Linux 7.2, provided that the hardware vendor has qualified it and provides full support for the particular HBA and storage array configuration. DIF/DIX is not supported on other configurations, it is not supported for use on the boot device, and it is not supported on virtualized guests.

At the current time, the following vendors are known to provide this support.

FUJITSU supports DIF and DIX on:

* EMULEX 16G FC HBA:

* EMULEX LPe16000/LPe16002, 10.2.254.0 BIOS, 10.4.255.23 FW, with:

* FUJITSU ETERNUS DX100 S3, DX200 S3, DX500 S3, DX600 S3, DX8100 S3, DX8700 S3, DX8900 S3, DX200F, DX60 S3

* QLOGIC 16G FC HBA:

* QLOGIC QLE2670/QLE2672, 3.28 BIOS, 8.00.00 FW, with:

* FUJITSU ETERNUS DX100 S3, DX200 S3, DX500 S3, DX600 S3, DX8100 S3, DX8700 S3, DX8900 S3, DX200F, DX60 S3

Note that T10 DIX requires database or some other software that provides generation and verification of checksums on disk blocks. No currently supported Linux file systems have this capability.

EMC supports DIF on:

* EMULEX 8G FC HBA:

* LPe12000-E and LPe12002-E with firmware 2.01a10 or later, with:

* EMC VMAX3 Series with Enginuity 5977; EMC Symmetrix VMAX Series with Enginuity 5876.82.57 and later

* EMULEX 16G FC HBA:

* LPe16000B-E and LPe16002B-E with firmware 10.0.803.25 or later, with:

* EMC VMAX3 Series with Enginuity 5977; EMC Symmetrix VMAX Series with Enginuity 5876.82.57 and later

* QLOGIC 16G FC HBA:

* QLE2670-E-SP and QLE2672-E-SP, with:

* EMC VMAX3 Series with Enginuity 5977; EMC Symmetrix VMAX Series with Enginuity 5876.82.57 and later

Please refer to the hardware vendor's support information for the latest status.

Support for DIF/DIX remains in Tech Preview for other HBAs and storage arrays.

Chapter 3. Authentication and Interoperability

Identity Management sets up a one-way trust by default

The **ipa trust-add** command now configures a one-way trust by default. One-way trusts enable users and groups in Active Directory (AD) to access resources in Identity Management (IdM) but not the other way around. Previously, the default trust configured by running **ipa trust-add** was a two-way trust.

IdM still allows the administrator to set up a two-way trust by adding the **--two-way=true** option to **ipa trust-add**.

openldap rebase to version 2.4.40

The *openldap* packages have been upgraded to upstream version 2.4.40, which provides a number of bug fixes and one enhancement over the previous version. Notably, the ORDERING matching rules have been added to the **ppolicy** attribute type descriptions. Among the fixed bugs are: The server no longer terminates unexpectedly when processing SRV records, and missing **objectClass** information has been added, which enables the user to modify the front-end configuration by standard means.

Cache authentication in SSSD

Authentication against cache without a reconnection attempt is now available in SSSD even in online mode. Authenticating directly against the network server repeatedly could cause excessive application latency, which could make the login process overly time-consuming.

SSSD enables UID and GID mapping on individual clients

It is now possible to map users to a different UID and GID on specific Red Hat Enterprise Linux clients through client-side configuration by using SSSD. This client-side override possibility can resolve problems caused by UID and GID duplication or ease transition from a legacy system that previously used different ID mapping.

Note that the overrides are stored in the SSSD cache; removing the cache therefore also removes the overrides.

SSSD can now deny SSH access to locked accounts

Previously, when SSSD used OpenLDAP as its authentication database, users could authenticate into the system successfully with an SSH key even after the user account was locked. The **ldap_access_order** parameter now accepts the **ppolicy** value, which can deny SSH access to the user in the described situation. For more information about using **ppolicy**, see the **ldap_access_order** description in the sssd-ldap(5) manual page.

The sudo utility now capable of verifying command checksum

The configuration of the sudo utility can now store the checksum of a command or script that is being permitted. When the command or script is run again, the checksum is compared to the stored checksum to verify that nothing has changed. If the command or binary is modified, the sudo utility refuses to run the command or logs a warning. This functionality makes it possible to correctly devolve responsibility and problem-solving activities if an incident occurs.

SSSD smart card support

SSSD now supports smart cards for local authentication. With this feature, the user can use a smart card to log on to the system using a text-based or graphical console, as well as local services such as the sudo service. The user places the smart card into the reader and provides the user name and the smart card PIN at the login prompt. If the certificate on the smart card is verified, the user is successfully authenticated.

Note that SSSD does not currently enable the user to acquire a Kerberos ticket using a smart card. To obtain a Kerberos ticket, the user is still required to authenticate using the kinit utility.

Support for multiple certificate profiles and user certificates

Identity Management now supports multiple profiles for issuing server and other certificates instead of only supporting a single server certificate profile. The profiles are stored in the Directory Server and shared between IdM replicas.

In addition, the administrator can now issue certificates to individual users. Previously, it was only possible to issue certificates to hosts and services.

Password Vault

A new feature to allow secure central storage of private user information, such as passwords and keys has been added to Identity Management. Password Vault is built on top of the Public Key Infrastructure (PKI) Key Recovery Authority (KRA) subsystem.

Kerberos HTTPS proxy in Identity Management

A Key Distribution Center (KDC) proxy function, interoperable with the Microsoft Kerberos KDC Proxy Protocol (MS-KKDCP) implementation, is now available in Identity Management and allows clients to access the KDC and **kpasswd** services by using HTTPS. System administrators can now expose the proxy at their network edge by a simple HTTPS reverse proxy without the need to set up and manage a dedicated application.

Background refresh of cached entries

SSSD now allows cached entries to be updated out-of-band in the background. Prior to this update, when the validity of cached entries expired, SSSD fetched them from the remote server and stored them in the database anew, which could be time consuming. With this update, entries are returned instantly because the back end keeps them updated at all times. Note that this causes a higher load on the server because SSSD downloads the entries periodically instead of only upon request.

Caching for `initgroups` operations

The SSSD fast memory cache now supports the **initgroups** operations, which enhances the speed of initgroups processing and improves the performance of some applications, for example GlusterFS and **slapi-nis**.

Negotiate authentication streamlined with `mod_auth_gssapi`

Identity Management now uses the **mod_auth_gssapi** module, which uses GSSAPI calls instead of direct Kerberos calls used by the previously used **mod_auth_kerb** module.

User life-cycle management capabilities

The user life-cycle management gives the administrator a greater degree of control over activating and deactivating user accounts. The administrator can now provision new user accounts by adding them to a stage area without fully activating them, activate inactive user accounts to make them fully operational, or deactivate user accounts without completely deleting them from the database.

User life-cycle management capabilities bring significant benefits to large IdM deployments. Note that users can be added to the stage area also directly from a standard LDAP client, using direct LDAP operations. Previously, IdM only supported managing users using IdM command-line tools or the IdM web UI.

SCEP support in `certmonger`

The **`certmonger`** service has been updated to support the Simple Certificate Enrollment Protocol (SCEP). It is now possible to issue a new certificate and renew or replace existing certificates over SCEP.

Apache modules for IdM now fully supported

The following Apache modules for Identity Management (IdM), added as Technology Preview in Red Hat Enterprise Linux 7.1, are now fully supported: **`mod_authnz_pam`**, **`mod_lookup_identity`**, and **`mod_intercept_form_submit`**. The Apache modules can be used by external applications to achieve tighter interaction with IdM beyond simple authentication.

NSS raises minimum accepted key strength values

The Network Security Services (NSS) library in Red Hat Enterprise Linux 7.2 no longer accepts Diffie-Hellman (DH) key exchange parameters smaller than 768 bits, nor RSA and DSA certificates with key sizes less than 1023 bits. Raising the minimum accepted key strength values prevents attacks exploiting known security vulnerabilities such as Logjam (CVE-2015-4000) and FREAK (CVE-2015-0204).

Note that attempts to connect to a server by using keys weaker than the new minimum values now fail, even though such connections worked in previous versions of Red Hat Enterprise Linux.

NSS enables TLS version 1.1 and 1.2 by default

Applications using protocol versions that NSS enables by default now additionally support the TLS version 1.1 and TLS version 1.2 protocols.

ECDSA certificates are now supported

Applications that use the default NSS cipher list now support connections to servers that use Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

OpenLDAP automatically chooses the NSS default cipher suites

OpenLDAP clients now automatically choose the Network Security Services (NSS) default cipher suites for communication with the server. It is no longer necessary to maintain the default cipher suites manually in the OpenLDAP source code.

Configuring an IdM server to be a trust agent now supported

Identity Management (IdM) distinguishes two types of IdM master servers: trust controllers and trust agents. Trust controllers run all the services required for establishing and maintaining a trust; trust agents only run services required to provide resolution of users and groups from trusted Active Directory forests to IdM clients enrolled with these IdM servers.

By default, running the **ipa-adtrust-install** command sets up the IdM server as a trust controller. To configure another IdM server to be a trust agent, pass the **--add-agents** option to **ipa-adtrust-install**.

Automated migration from WinSync to trusts now supported

The new **ipa-winsync-migrate** utility enables seamless migration from synchronization-based integration using WinSync to integration based on Active Directory (AD) trust. The utility automatically migrates all users synchronized using WinSync from a specified AD forest. Previously, migration from synchronization to trust could only be performed manually using ID views.

For more information about **ipa-winsync-migrate**, see the ipa-winsync-migrate(1) man page.

Multi-step prompting for one-time and long-term passwords

When using a one-time password (a token) together with a long-term password to log in, the user is prompted for both passwords separately. This results in better user experience when using one-time passwords as well as a safer long-term password extraction, which allows long-term password caching to be used for offline authentication.

LPK schema for OpenLDAP now available in the LDIF format

LDIF is the new default format for the OpenLDAP import schema, and the *openssh-ldap* package now provides the LDAP Public Key (LPK) schema in the LDIF format as well. Therefore, administrators can directly import the LDIF schema when setting up public-key authentication based on LDAP.

Cyrus can authenticate to AD and IdM servers again

An upstream release of the *cyrus-sasl* packages introduced a non-backward compatible change that prevented Cyrus from authenticating against older SASL implementations. Consequently, Red Hat Enterprise Linux 7 was not able to authenticate to Active Directory (AD) and Red Hat Enterprise Linux 6 Identity Management (IdM) servers. The upstream change has been reverted and Cyrus can now authenticate to AD and IdM servers as expected.

SSSD supports overriding automatically discovered AD site

The Active Directory (AD) DNS site to which the client connects is discovered automatically by default. However, the default automatic search might not discover the most suitable AD site in certain setups. In such situations, you can now define the DNS site manually using the **ad_site** parameter in the **[domain/NAME]** section of the **/etc/sssd/sssd.conf** file.

Support for SAML ECP has been added

The lasso packages have been rebased to version 2.5.0 and the mod_auth_mellon packages have been rebased to version 0.11.0 in order to add support for Security Assertion Markup Language (SAML) Enhanced Client or Proxy (ECP). SAML ECP is an alternative SAML profile that allows non-browser-based Single Sign On (SSO).

Chapter 4. Clustering

systemd and pacemaker now coordinate correctly during system shutdown

Previously, systemd and pacemaker did not coordinate correctly during system shutdown, which caused pacemaker resources not to be terminated properly. With this update, pacemaker is ordered to stop before dbus and other systemd services that pacemaker started. This allows both pacemaker and the resources that pacemaker manages to shut down properly.

The `pcs resource move` and `pcs resource ban` commands now display a warning message to clarify the commands' behavior

The **`pcs resource move`** command and the **`pcs resource ban`** commands create location constraints that that effectively ban the resource from running on the current node until the constraint is removed or until the constraint lifetime expires. This behavior had previously not been clear to users. These commands now display a warning message explaining this behavior, and the help screens and documentation for these commands have been clarified.

New command to move a Pacemaker resource to its preferred node

After a Pacemaker resource has moved, either due to a failover or to an administrator manually moving the node, it will not necessarily move back to its original node even after the circumstances that caused the failover have been corrected. You can now use the **`pcs resource relocate run`** command to move a resource to its preferred node, as determined by current cluster status, constraints, location of resources and other settings. You can also use the **`pcs resource relocate show`** command to display migrated resources. For information on these commands, see the High Availability Add-On Reference.

Support for clufter command for transforming and analyzing cluster configuration formats

The clufter command provides a tool for transforming and analyzing cluster configuration formats. The clufter command can be used to assist with migration from an older stack configuration to a newer configuration that leverages Pacemaker. For information on the capabilities of the clufter command, see the clufter(1) man page or the output of the **`clufter -h`** command.

Simplified method for configuring fencing for redundant power supplies in a cluster

When configuring fencing for redundant power supplies, you must ensure that when the power supplies are rebooted both power supplies are turned off before either power supply is turned back on. If the node never completely loses power, the node may not release its resources. This opens up the possibility of nodes accessing these resources simultaneously and corrupting them.

Prior to Red Hat Enterprise Linux 7.2, you needed to explicitly configure different versions of the devices which used either the 'on' or 'off' actions. Since Red Hat Enterprise Linux 7.2, it is now only required to define each device once and to specify that both are required to fence the node.

For information on configuring fencing for redundant power supplies, see the **`Fencing: Configuring STONITH`** chapter of the High Availability Add-On Reference manual.

New --port-as-ip option for fencing agents

Fence agents used only with single devices required complex configuration in pacemaker. It is now possible to use the **`--port-as-ip`** option to enter the IP address in the **`port`** option.

Chapter 5. Compiler and Tools

tail --follow now works properly on files on Veritas Clustered file system (VXFS)

Veritas Clustered file system (VXFS) is a remote file system, and for remote file systems, **tail** cannot use **inotify** functionality for **--follow** mode. Veritas Clustered file system has now been added to the list of remote file systems, for which polling mode is used instead of **inotify**. **tail --follow** can now work properly even when used on files on VXFS.

The dd command now capable of showing transfer progress

The **dd** command, which is used for copying files by bytes, now provides the **status=progress** option for showing the progress of the transfer. This is especially useful for transfers of large files because it allows the user to estimate the time left and detect potential issues with the transfer.

Improved wait times in libcurl

The **libcurl** library used an unnecessarily long blocking delay for actions with no active file descriptors, even for short operations. This meant that some actions, such as resolving a host name using **/etc/hosts**, took an artificially long time to complete. The blocking code in **libcurl** has now been modified so that the initial delay is short, and gradually increases until an event occurs. Fast **libcurl** operations now complete more quickly.

The libcurl library now implements a non-blocking SSL handshake

Previously, the **libcurl** library did not implement a non-blocking SSL handshake, which negatively affected performance of applications based on the **libcurl** multi API. To solve this problem, the non-blocking SSL handshake has been implemented in **libcurl**, and the **libcurl** multi API now immediately returns the control back to the application whenever it cannot read or write data from or to the underlying network socket.

GDB on IBM Power Systems no longer fails when accessing the symbol table

Previously, GDB on 64-bit IBM Power Systems incorrectly deallocated an important variable that held the symbol table for the binary being debugged, which caused a segmentation fault when GDB tried to access that symbol table. To solve this issue, this specific variable has been made persistent, and GDB can now access the needed information later during the debugging session, without reading an invalid memory region.

nscd updated to automatically reload configuration data

This update of Name Server Caching Daemon (nscd) adds a system of inotify-based monitoring and stat-based backup monitoring for nscd configuration files, so that nscd now correctly detects changes to its configuration and reloads the data. This prevents nscd from returning stale data.

The dlopen library function no longer crashes on recursive calls

Previously, a defect in the library function **dlopen** could cause recursive calls to this function to crash or abort with a library assertion. Recursive calls are possible if a user-provided **malloc** implementation calls **dlopen**.

The implementation is now reentrant, and recursive calls no longer crash or abort with an assertion.

The operf tool now recognizes static huge page identifiers

Previously, when profiling performance of Java just-in-time (JIT) compiled code with static huge pages enabled, OProfile's operf command recorded a large number of event samples to anonymous memory (in anon_hugepage) instead of to the appropriate Java method. With this update, operf recognizes the static huge page identifiers and correctly maps samples to Java methods when using statically allocated huge pages.

`rsync -X` now works correctly

Previously, the **rsync** tool changed the file ownership after, not before, setting security attributes. As a consequence, the security attributes on the target were missing, and running the **`rsync -X`** command did not work correctly under certain circumstances. With this update, the order of the operations has been switched, and **rsync** now changes the ownerships before setting the security attributes. As a result, the security attributes are present as expected in the described situation.

Subversion executables now built with full RELRO data

The executables supplied with the *subversion* package are now built with fully read-only relocation data (RELRO), which provides protection against some types of memory corruption attacks. As a result, it will be harder to successfully exploit Subversion if future vulnerabilities are discovered.

The thread extension in TCL now works correctly

Previously, the threading support in the Tool Command Language (TCL) was not implemented optimally. If the fork() call was used together with thread extension enabled in the TCL interpreter, the process could become unresponsive. Because of that, the TCL interpreter and the TK application were previously shipped with the thread extension disabled. As a consequence, third-party applications depending on threaded TCL or TK did not work correctly. A patch has been implemented to fix this bug, and TCL and TK now have the thread extension enabled by default.

AES cipher suites can be explicitly enabled or disabled for TLS

With the updated curl packages, it is possible to explicitly enable or disable new Advanced Encryption Standard (AES) cipher suites to be used for the TLS protocol.

OpenJDK 7 now supports ECC

With this update, OpenJDK 7 supports Elliptic Curve Cryptography (ECC) and the associated ciphers for TLS connections. ECC is in most cases preferable to older cryptographic solutions for making secure network connections.

ABRT is now able to save a core_backtrace file instead of a whole coredump

ABRT can now be configured to generate a backtrace without writing a coredump file to the disk. This can save time when working with processes which have allocated large blocks of memory. This feature can be enabled by setting the **`CreateCoreBacktrace`** option to **`yes`** and the **`SaveFullCore`** option to **`no`** in the **`/etc/abrt/plugins/CCpp.conf`** file.

Security features added to the Python standard library

A number of security enhancements, which are described in the 466 Python Enhancent Proposal (http://legacy.python.org/dev/peps/pep-0466/), have been backported to the Python standard library. The security enhancements include, for example, new features in the **`ssl`** module, such as support for Server Name Indication (SNI) as well as support for new TLSv1.x protocols, new hash algorithms in the **`hashlib module`**, and much more.

New global settings for SSL/TLS certificate verification in the Python standard library

New options have been added that allow users to globally enable or disable SSL/TLS certificate verification in the HTTP clients (such as urllib, httplib, or xmlrpclib) of the Python standard library. The options are described in the 493 Python Enhancement Proposal (https://www.python.org/dev/peps/pep-0493/). The default value is to not verify certificates. For details, see https://access.redhat.com/articles/2039753.

Chapter 6. Desktop

GNOME rebase to version 3.14

The **GNOME Desktop** has been upgraded to upstream version 3.14 (with some minor additions from 3.16), which includes new features and a number of enhancements. Namely:

Red Hat Enterprise Linux 7.2 adds **GNOME Software**, a new way to install and manage software on the user's system based on a yum backend. GNOME PackageKit remains to be the default updater for GNOME (also installed by default). With **GNOME Software**, the user manages an integrated place for software related tasks, such as browsing, installing and removing applications, and viewing and installing software updates.

On the Top Bar, the newly-named **System Status Menu** groups together all of the indicators and applets otherwise accessed individually – brightness slider, improved airplane mode, connecting to Wi-Fi networks, Bluetooth, Volume, and so on – into one coherent and compact menu. Regarding Wi-Fi, GNOME 3.14 provides improved support for Wi-Fi hotspots. When connecting to a Wi-Fi portal that requires authentication, GNOME now automatically shows the login page as a part of the connection process.

The new design of the **gedit** text editor incorporates all of previous features into a more compact interface, which gives more space for work. Use of popovers for selecting the document format and tab width is more efficient compared to the previous use of dialogs and menus. Consolidated sidebar controls also give more space for content while retaining the original functionality. Other notable improvements include new shortcuts for opening the last closed tab with **Ctrl+Shift+T** and for changing case.

Nautilus, the GNOME file manager, now uses the **Shift+Ctrl+Z** key combination, not **Ctrl+Y**, for the **redo** operation. Also, a header bar, instead of a toolbar, is now used.

GNOME 3.14 includes a reimagined **Videos** application. Modern in style, the new version allows the user to browse videos on the computer as well as online video channels. **Videos** also includes a redesigned playback view. This provides a more streamlined experience than the earlier version: floating playback controls hide when the user does not need them, and the fullscreen playback view also has a new more refined look.

Evince features improved accessibility for reading PDF files. The new version of the document viewer uses a header bar to give more space to your documents. When it is launched without a document being specified, **Evince** also shows a useful overview of your recent documents. The latest **Evince** version also includes **High Resolution Display Support** and enhanced accessibility, with links, images and form fields all being available from assistive technologies.

The new version of GNOME **Weather** application makes use of GNOME's new geolocation framework to automatically show the weather for your current location, and a new layout provides an effective way to read weather forecasts.

This release also brings improved support for comments in **LibreOffice** - import and export of nested comments in the ODF, DOC, DOCX and RTF filters, printing comments in margins, and formatting all comments.

Also, redesigned GNOME provides support for browsing more pictures in **Photos**, and improved touchscreen support, with multi-touch gestures for both the system and applications.

The GNOME application for virtual and remote machines, **Boxes**, introduces snapshots. **Boxes** now provide automatic downloading, running multiple boxes in separate windows, and user interface improvements, including improved fullscreen behavior and thumbnails.

The GNOME **Help** documentation viewer has been redesigned to be consistent with other GNOME 3 applications. Help now uses a header bar, has an integrated search function, and bookmarking interface.

GTK+ 3.14 includes a number of bug fixes and enhancements, such as automatic loading of menus from resources, multi-selection support in **GtkListBox**, property bindings in **GtkBuilder** files, support for drawing outside a widget's allocation (gtk_widget_set_clip()), new transition types in GtkStack, and file loading and saving with **GtkSourceView**. In addition, **GTK+** now provides support for gesture interaction. With 3.14, the majority of common multitouch gestures are available for use in GTK+ applications, such as tap, drag, swipe, pinch, and rotate. Gestures can be added to existing GTK+ applications using **GtkGesture**.

A GNOME Shell Extension, **Looking Glass Inspector**, has obtained a number of features for developers: showing all methods, classes, and so on in a namespace upon inspection, object inspector history expansion, or copying **Looking Glass** results as strings, and passing through events to gnome-shell.

The **High Resolution Display Support** feature has been extended to include all the key aspects of the core GNOME 3 experience, including the Activities Overview, animations in the Activities Overview along with new window animations, Top Bar, lock screen and system dialogs.

As far as GNOME Extensions are concerned, this release introduces support for alternative dock positioning, including the bottom side of the screen, in **Simple Dock**, a dock for the Gnome Shell.

The ibus-gtk2 package now updates the immodules.cache file

Previously, the **update-gtk-immodules** script searched for a no longer existing **/etc/gtk-2.0/$host** directory. Consequently, the post-installation script of the *ibus-gtk2* package failed and exited without creating or updating the cache. The post-installation script has been changed to replace **update-gtk-immodules** with **gtk-query-immodules-2.0-BITS**, and the problem no longer occurs.

Chapter 7. File Systems

gfs2-utils rebase to version 3.1.8

The *gfs2-utils* package has been rebased to version 3.1.8, which provides important fixes and a number of enhancements:

* The performance of the **`fsck.gfs2`**, **`mkfs.gfs2`**, and **`gfs2_edit`** utilities has been improved.

* The **`fsck.gfs2`** utility now performs better checking of journals, the jindex, system inodes, and the inode 'goal' values.

* The **`gfs2_jadd`** and **`gfs2_grow`** utilities are now separate programs instead of symlinks to **`mkfs.gfs2`**.

* The test suite and related documentation have been improved.

* The package no longer depends on Perl.

GFS2 now prevents users from exceeding their quotas

Previously, GFS2 only checked quota violations after the completion of operations, which could result in users or groups exceeding their allotted quotas. This behavior has been fixed, and GFS2 now predicts how many blocks an operation would allocate and checks if allocating them would violate quotas. Operations that would result in quota violations are disallowed, and users thus never exceed their allotted quotas.

XFS rebase to version 4.1

XFS has been upgraded to upstream version 4.1 including minor bug fixes, refactorings, reworks of certain internal mechanisms, such as logging, pcpu accounting, and new mmap locking. On top of the upstream changes, this update extends the rename() function to add cross-rename (a symmetric variant of rename()) and whiteout handling.

cifs rebase to version 3.17

The CIFS module has been upgraded to upstream version 3.17, which provides various minor fixes and new features for Server Message Block 2 and 3 (SMB2 and SMB3).

Changes in NFS in Red Hat Enterprise Linux 7.2

Fallocate support allows preallocation of files on the server. The SEEK_HOLE and SEEK_DATA extensions to the fseek() function make it possible to locate holes or data quickly and efficiently. Red Hat Enterprise Linux 7.2 also adds support for flexible file layout on NFSv4 clients described in the Technology Previews section.

Chapter 8. Hardware Enablement

OSA-Express5s cards support in qethqoat

Support for OSA-Express5s cards was added to the qethqoat tool, part of the s390utils package, in Red Hat Enterprise Linux 7.1 as a Technology Preview. This enhancement update provides full support of the extended serviceability of network and card setups for OSA-Express5s cards.

Chapter 9. Installation and Booting

Fixed network setup in initrd if network configuration is provided in Kickstart

Previously, the installer was failing to set up or reconfigure network interfaces in **initrd**, if these interfaces were defined in Kickstart files. This could cause the installation to fail and enter emergency mode if network access was required by other commands in the Kickstart file.

This issue is now resolved and Anaconda now properly handles network configuration from Kickstart files in **initrd**, early in the boot process.

Anaconda now supports creating cached logical volumes

The installer now supports creating cached LVM logical volumes and installing the system onto those volumes.

Currently, this approach is only supported in Kickstart. To create a cached logical volume, use the new **--cachepvs=**, **--cachesize=**, and **--cachemode=** options of the **logvol** Kickstart command.

See the Red Hat Enterprise Linux 7 Installation Guide for detailed information about these new options.

Improved sorting of GRUB2 boot menu

An issue with the sorting mechanism used by the **grub2-mkconfig** command could cause the grub.cfg configuration file to be generated with available kernels sorted incorrectly.

GRUB2 now uses the *rpmdevtools* package to sort available kernels and the configuration file is being generated correctly with the most recent kernel version listed at the top.

Anaconda now properly reverts disk actions when disk selection changes

Previously, Anaconda and Blivet did not properly revert actions scheduled on disks when disk selection changed, causing various issues. With this update, Anaconda has been fixed to create a snapshot of the original storage configuration and return to it when disk selection changes, thus completely reverting all actions scheduled for disks.

Improved detection of device-mapper disk names

In the previous release of Red Hat Enterprise Linux 7, it was possible for the installer to crash when installing on disks which previously contained LVM logical volumes and the metadata for those volumes was still present. The installer could not recognize correct **device-mapper** names and the process of creating new LVM logical volumes would fail.

The method used to obtain **device-mapper** device names has been updated and installation on disks which contain existing LVM metadata is now more reliable.

Fixed handling of PReP Boot during partitioning

In some circumstances, the **PReP Boot** partition on IBM Power Systems could be set to an invalid size during custom partitioning. In that situation, removing any partition caused the installer to crash.

Checks are now implemented in *anaconda* to ensure that the partition is always sized correctly between **4096 KiB** and **10 MiB**. Additionally, it is no longer necessary to change the format of the **PReP Boot** partition in order to change its size.

EFI partitions on RAID1 devices

EFI System Partitions may now be created on a RAID1 device, this is to enable system recovery when one boot disk fails. However, because the system is only guaranteed to discover one EFI System Partition, if the volume of the ESP that is discovered by the firmware becomes corrupt (but still appears to be a valid ESP), and both **Boot####** and **BootOrder** also become corrupt, then the boot order will not be rebuilt automatically. In this case, the system should still boot manually from the second disk.

Text mode installation no longer crashes during network configuration

Previously, in the Network Configuration screen in the interactive text mode installer, using a space when specifying nameservers caused the installer to crash.

Anaconda now handles spaces in nameserver definitions in text mode correctly and the installer no longer crashes if a space is used to separate nameserver addresses.

Rescue mode screens on IBM System z are no longer cut off

Previously, the second and third screen in rescue mode on IBM System z servers were being displayed improperly and parts of the interface were cut off. Rescue mode on this architecture has been improved and all screens now function correctly.

OpenSCAP add-on in Anaconda

It is now possible to apply Security Content Automation Protocol (SCAP) content during the installation process. This new installer add-on provides a reliable and easy way to configure a security policy without having to rely on custom scripts.

This add-on provides a new Kickstart section ("%addon org_fedora_oscap") as well as a new screen in the graphical user interface during an interactive installation. All three parts are documented in the Red Hat Enterprise Linux 7 Installation Guide.

Applying a security policy during installation will perform various changes during and immediately after the installation, depending on which policy you enable. If a profile is selected, the *openscap-scanner* package (an OpenSCAP compliance scanning tool) is added to your package selection and an initial compliance scan is performed after the installation finishes. Results of this scan are saved into **/root/openscap_data**.

Several profiles are provided on installation media by the *scap-security-guide* package. You can also load other content as a datastream, archive, or an RPM package from an HTTP, HTTPS or FTP server if needed.

Note that applying a security policy is not necessary on all systems. This add-on should only be used when a specific policy is mandated by your organization's rules or government regulations, otherwise the add-on can be left in its default state which does not apply any security policy.

Anaconda no longer times out when waiting for a Kickstart file on a CD or DVD

Previously, if Anaconda was configured to load a Kickstart file from optical media using the **`inst.ks=cdrom:/ks.cfg`** command, and the system was also booted from a CD or DVD, the installer only waited 30 seconds for the user to swap the disk. After this time window passed, the system entered emergency mode.

With this update, Anaconda has been modified to never time out when waiting for the user to provide a Kickstart file on a CD or DVD. If the **`inst.ks=cdrom`** boot options is used and the Kickstart file is not detected, Anaconda displays a prompt and waits until the user provides the file or reboots.

Chapter 10. Kernel

The SHMMAX and SHMALL kernel parameters returned to default values

Previously, the values of the **kernel.shmmax** and **kernel.shmall** parameters, which were set in the **/usr/lib/sysctl.d/00-system.conf** file, were too low. As a consequence, some applications, such as SAP, could not function properly. The unsuitable overrides have been removed, and the kernel defaults, which are sufficiently high, are now used.

Transparent huge pages no longer cause memory corruption

Transparent huge pages were not being correctly synchronized during read and write operations. In some circumstances, this resulted in memory corruption when transparent huge pages were enabled. Memory barriers have been added to transparent huge page handling so that this memory corruption no longer occurs.

SCSI LIO rebase

The SCSI kernel target, LIO, has been rebased from Linux-4.0.stable. This includes many bug fixes, most critically for iSER, but also includes added support for the XCOPY, WRITE SAME, and ATS commands; and DIF data integrity support.

makedumpfile now supports the new sadump format representing up to 16 TB of physical memory

The makedumpfile command now supports the new sadump format that can represent more than 16 TB of physical memory space. This allows users of makedumpfile to read dump files over 16 TB, generated by sadump on certain upcoming server models.

Removing or upgrading kernel no longer displays a warning

The weak-modules script, which is used by **kmod** to manage kABI-compatible module symbolic links, was previously removing the **/lib/modules/<version>/weak-updates** directory when removing files associated with a kernel. This directory is owned by the *kernel* package and removing it caused inconsistency between the file system and the state expected by **rpm**. This caused a warning to be displayed every time a kernel was upgraded or removed.

The script has been updated to remove the contents of the **weak-updates** directory but leave the directory itself, and warnings are no longer being displayed.

New package: libevdev

Libevdev is a low-level library for the Linux kernel input event device interface. It provides safe interfaces to query device capabilities and process events from devices. Current versions of xorg-x11-drv-evdev and xorg-x11-drv-synaptics require this library as a dependency.

Tuned can now run in no-daemon mode

Previously, Tuned could run only as a daemon, which could affect performance of small systems because of the memory footprint of the Tuned daemon. With this update, a no-daemon (one shot) mode, which does not require any resident memory, has been added into Tuned. The no-daemon mode is disabled by default because much of Tuned functionality is missing in this mode.

New package: tuned-profiles-realtime

The *tuned-profiles-realtime* package has been added to Red Hat Enterprise Linux Server and Red Hat Enterprise Linux for Real Time. It contains a realtime profile used by the **tuned** utility to perform CPU isolation and IRQ tuning. When the profile is activated, it reads its variable section, which specifies the CPUs to be isolated, and moves all threads that may be moved off those CPU cores.

SCSI error messages can now be interpreted comfortably

Previous kernel changes to the printk() function had resulted in Small Computer System Interface (SCSI) error messages being logged across multiple lines. As a consequence, if multiple errors occurred across different devices, it could become difficult to interpret the error messages correctly. This update changes the SCSI error logging code to log error messages using the dev_printk() option, which associates each error message with the device that generated the error.

libATA subsystem and drivers updated

This enhancement update provides number of bug fixes and enhancements of the libATA subsystem and drivers.

FCoE and DCB have been upgraded

Fibre Channel over Ethernet (FCoE) and Data Center Bridging (DCB) kernel components have been upgraded to the latest upstream versions, which provides a number of bug fixes and enhancements over the previous versions.

`perf` rebase to version 4.1

The *perf* packages have been upgraded to upstream version 4.1, which provides a number of performance and stability fixes and enhancements over the previous version. Notably, this rebase adds the Intel Cache QoS Monitoring and AMD IBS Ops features and provides support for Intel Xeon v4 for compressed kernel modules, for parametrized events and support to specify breakpoint length. In addition, a number of options have been added to the **perf** tool, such as the **`--system-wide`**, **`top -z`**, **`top -w`**, **`trace --filter-pids`**, and **`trace --event`** options.

Support for TPM 2.0

This update adds driver-level support for version 2.0 compliant Trusted Platform Module (TPM) devices.

`turbostat` now provides correct output

Previously, the **`turbostat`** tool detected if system had the MSR device support by reading the /dev/cpu/0/msr file for **`cpu0`** instead of **`cpu`**. As a consequence, disabling a CPU caused the CPUs to be deleted from turbostat output. This bug has been fixed, and running the **`turbostat ls`** command now returns correct output.

`turbostat` now supports Intel Xeon v5 processors

This enhancement adds Intel Xeon v5 processor support to the **`turbostat`** tool.

the `zswap` tool makes use of the `zpool` API

Previously, the **zswap** tool directly used **zbud**, a storage pool that stores compressed pages at a ratio of 2:1 (when full). This update introduces the **zpool** API that provides access to either the **zbud** or **zsmalloc** pools: **zsmalloc** stores compressed pages at a potential higher density, resulting in more reclaimed memory for highly compressible pages. Within this update, **zsmalloc** has been promoted to the /mm drivers so that **zpool** works as intended.

The /proc/pid/cmdline file length is now unlimited

The /proc/pid/cmdline file length limit for the **ps** command was previously hard-coded in the kernel to 4096 characters. This update makes sure the length of /proc/pid/cmdline is unlimited, which is especially useful for listing processes with long command line arguments.

Support for dma_rmb and dma_wmb now provided

This update introduces two new primitives for synchronizing cache coherent memory writes and reads, dma_wmb() and dma_rmb(). This feature will be available for appropriate use in drivers.

qib HCA driver connection

Due to a mismatch in SRP LOGIN ID, the SRP target previously failed to connect over the qib HCA device driver. This update fixes the bug, and the aforementioned connection can now be established successfully.

Increase in memory limit

Staring with Red Hat Enterprise Linux 7.2, maximum supported memory limit on AMD64 and Intel 64 systems has been increased from 6 TB to 12 TB.

Chapter 11. Networking

SNMP now correctly obeys the clientaddr directive over IPv6

Previously, the **clientaddr** option in **snmp.conf** only affected outgoing messages sent over IPv4. With this release, the outgoing IPv6 messages are correctly sent from the interface specified by **clientaddr**.

tcpdump supports -J, -j, and --time-stamp-precision options

As **kernel**, **glibc**, and **libpcap** now provide APIs to obtain nanosecond resolutions time stamps, **tcpdump** has been updated to leverage this functionality. Users can now query which time stamp sources are available (-J), set a specific time stamp source (-j), and request time stamps with a specified resolution (--time-stamp-precision).

TCP/IP rebase to version 3.18

TCP/IP stack has been upgraded to upstream version 3.18, which provides a number of bug fixes and enhancements over the previous version. Notably, this update fixes TCP fast open extension, which now works as expected when using IPv6. In addition, this update provides support for optional TCP autocorking and implements Data Center TCP (DCTCP).

NetworkManager libreswan rebase to version 1.0.6

A number of bug fixes and enhancements have been incorporated from upstream, for example:

* Password handling is now more robust

* Connection start and stop is now more robust

* Default routing is now autodetected from pushed routes

* Added support for interactive password requests

* Fixed erroneous import and export capability advertisement.

NetworkManager now supports setting the MTU of a bonded interface

Both 'nmcli' and the GUI interface now allow the setting of MTU on a bonded interface.

NetworkManager now validates IPv6 Router Advertisement MTU options before applying them

Malicious or misconfigured nodes could send an IPv6 MTU that would make further network communication problematic or impossible if applied. NetworkManager now gracefully handles these events and maintains IPv6 connectivity.

IPv6 Privacy extensions now enabled by default

To determine and set IPv6 privacy settings at device activation, NetworkManager now checks its network configuration in NetworkManager.conf by default, and falls back to **/proc/sys/net/ipv6/conf/default/use_tempaddr** if necessary.

The control-center Network Panel now displays WiFi device capabilities

Supported operating frequencies of WiFi devices are now displayed in the control-center network panel.

NetworkManager now gracefully handles route conflicts when multiple interfaces point to the same gateway

NetworkManager now keeps track of configured routes and avoids attempts to set conflicting routes. When a conflicting route is no longer active, it is removed.

Fix for network blackout with multihomed connections

NetworkManager now avoids a network blackout when activating the second device in a multihomed connection.

New option to prevent NetworkManager from overriding `ip route add`

The new 'never-default' option has been added to the connection IP configuration. This option prevents NetworkManager from setting the default route itself, allowing the administrator to set different default routes as required.

Fix for legacy network.service errors when Carrier Down is detected on some hardware

When a device has no carrier during boot, NetworkManager will wait for the carrier to be detected instead of causing activation to fail immediately.

NetworkManager now supports Wake On Lan

The nmcli utility now allows **`Wake on Lan`** to be set on a per device basis.

Improved support for firewalld zones with VPN connections

When a firewall zone is configured for a device-based VPN connection, the zone is now correctly configured in firewalld.

Fair Queue packet scheduler now supported

The Fair Queue packet scheduler, known as **`fq`**, has been added to Red Hat Enterprise Linux 7.2 and can be selected using the **`tc`** (traffic controller) utility.

Added support for transmit coalescing

The **`xmit_more`** extension has been implemented, improving transmit performance of virtio-net and other drivers, especially when TSO (TCP Segmentation Offload) is disabled.

Improved network frame receiving performance

By refactoring the code to eliminate IRQ save and restore operations in NAPI memory allocation, latency when receiving network frames has been reduced.

Significantly improved performance of route lookups

The IPv4 FIB (Forward Information Base) code has been updated from upstream to improve performance.

Network Namespace support for Virtual Interfaces

The netns id is now supported on virtual interfaces, allowing reliable tracking of linked network interfaces across network namespace boundaries.

Docker and LXC containers can now read net.ipv4.ip_local_port_range

Network name space support for the net.ipv4.ip_local_port_range sysctl has been added, improving container support for software that requires access to this information.

Improved reporting of autoconfigured IPv6 routes by the 'ip' tool

The **`ip`** tool could not get the mtu or hoplimit information from a Route Advertisement, this has been fixed.

Dual-stack socket options are now correctly exported

AF_INET6 sockets are only exclusive to IPv6 when IPV6_V6ONLY is set. In all other cases the socket is also IPv4 capable. This information is now properly exported and can be interrogated using iproute2.

Data Center TCP Now Supported

This release includes an implementation of DCTCP to improve network performance in Data Center environments. the parameter **`dctcp`** can be set either in **`sysctl`** or on a per route basis with **`ip route`**.

Per Route Congestion Control

To enable different congestion control algorithms on a per route basis, the **`congctl`** parameter has been added to **`ip route`**.

Improved Congestion Window handling for TCP Cubic and Reno when using GRO

The method to determine bandwidth and congestion window sizing has been improved, reducing the number of ACK packets required for transmission of large volumes of data.

TCP Pacing is now supported

The parameter **`SO_MAX_PACING_RATE`** has been added. This enables greater control of throughput rate for environment where this is a consideration.

Support for both client and server TFO

The TCP Fast Open feature has been added, using the RFC 7413 assigned option number.

Mitigation of TCP ACK loops

Handling of duplicated TCP ACKs has been improved, preventing some problems with buggy or potentially malicious middleboxes.

Minimal support for secondary endpoints with nf_conntrack_proto_sctp

Basic multihoming support has been added to SCTP.

AF_UNIX implementation rebased

The AF_UNIX (sometimes called AF_LOCAL) code has been updated to include many fixes and enhancements. In particular, sendpage and splice (also known as zerocopy) are now supported.

Kernel tunneling support rebased to upstream

The kernel tunneling drivers have been updated from kernel 4, bringing in many fixes and enhancements, especially for VXLAN.

Added support for crossing network namespaces to GRE

Both gre and ip6gre now have support for x-netns.

Improved performance when running Virtual Machine Traffic over VXLAN

The transmit flow hashing code has been updated, resulting in improved performance when traffic originating from a virtual machine is directed into a tunnel.

Improved offloading for VLAN frames received in a VXLAN or from GRE tunnels

A number of changes have been introduced to enable GRO support and improve performance under VXLAN and NVGRE tunneling.

Improved performance of Open vSwitch tunneling

The **tx-nocache-copy** device feature is now disabled by default. The previous default created a significant overhead for many workloads and particularly for OVS tunnels running over a VXLAN.

Improved IPsec Handling

IPsec has been updated to provide many fixes and some enhancements. Of particular note is that this release now provides the ability to match on outgoing interfaces.

Inclusion of VTI6 support including netns capabilities

Virtual Tunnel Interfaces for IPv6, including netns capabilities, have been added to the kernel.

Default value of nf_conntrack_buckets increased

If not specified as parameter during module loading, the default number of buckets is calculated through dividing total memory by 16384 to determine the number of buckets. The hash table will never have fewer than 32 and is limited to 16384 buckets. For systems with more than 4GB of memory however, this limit will be 65536 buckets.

Improvements in memory usage for iptables on large SMP machines

Previously, large iptables rulesets could use significant amounts of memory unnecessarily, this was due to storing the ruleset on a per (possible) CPU basis. The memory overhead has been reduced by changing the way rulesets are stored.

Network bonding driver updated

To improve maintainability, the kernel network bonding driver has been updated to bring it in line with upstream source.

Kernel netlink interfaces for bonding and 802.3ad (LACP)

Additional netlink interfaces for reading and setting bonding parameters on LACP devices have been added to the kernel.

Improvements in performance for mactap and macvtap with VLANs

Several low throughput issues involving segmentation problems have been addressed:

* Communicating with e1000 devices to virtio devices over mactap.

* Communicating with an external host when using VLANs in the guest.

* Communicating with the KVM host over a VLAN in both the guest and host.

Improved ethtool network querying

The network-querying capabilities of the ethtool utility were enhanced in a Technology Preview for Red Hat Enterprise Linux 7.1 on IBM System z and are fully supported as of Red Hat Enterprise Linux 7.2. As a result, when using hardware compatible with the improved querying, ethtool provides improved monitoring options, and displays network card settings and values more accurately.

Chapter 12. Security

GSSAPI key-exchange algorithms can now be selectively disabled

In view of the Logjam security vulnerability, the **`gss-group1-sha1-*`** key-exchange methods are no longer considered secure. While there was the possibility to disable this key-exchange method as a normal key exchange, it was not possible to disable it as a GSSAPI key exchange. With this update, the administrator can selectively disable this or other algorithms used by the GSSAPI key exchange.

SELinux policy for Red Hat Gluster Storage has been added

Previously, SELinux policy for Red Hat Gluster Storage (RHGS) components was missing, and Gluster worked correctly only when SELinux was in permissive mode. With this update, SELinux policy rules for the **`glusterd`** (glusterFS Management Service), **`glusterfsd`** (NFS sever), **`smbd`**, **`nfsd`**, **`rpcd`**, adn **`ctdbd`** processes have been updated providing SELinux support for Gluster.

openscap rebase to version 1.2.5

The *openscap* packages have been upgraded to upstream version 1.2.5, which provides a number of bug fixes and enhancements over the previous version.

Notable enhancements include:

* Support for OVAL version 5.11, which brings multiple improvements such as for systemd properties

* Introduced native support of **`xml.bz2`** input files

* Introduced the **`oscap-ssh`** tool for assessing remote systems

* Introduced the **`oscap-docker`** tool for assessing containers/images

scap-security-guide rebase to version 0.1.25

The *scap-security-guide* tool has been upgraded to upstream version 0.1.25, which provides a number of bug fixes and enhancements over the previous version.

Notable enhancements include:

* New security profiles for Red Hat Enterprise Linux 7 Server: Common Profile for General-Purpose Systems, Draft PCI-DSS v3 Control Baseline, Standard System Security Profile, and Draft STIG for Red Hat Enterprise Linux 7 Server.

* New security benchmarks for Firefox and Java Runtime Environment (JRE) components running on Red Hat Enterprise Linux 6 and 7.

* New **`scap-security-guide-doc`** subpackage, which contains HTML-formatted documents containing security guides generated from XCCDF benchmarks (for every security profile shipped in security benchmarks for Red Hat Enterprise Linux 6 and 7, Firefox, and JRE).

Chapter 13. Servers and Services

The ErrorPolicy directive is now validated

The ErrorPolicy configuration directive was not validated on startup, and an unintended default error policy could be used without a warning. The directive is now validated on startup and reset to the default if the configured value is incorrect. The intended policy is used, or a warning message is logged.

CUPS now disables SSLv3 encryption by default

Previously, it was not possible to disable SSLv3 encryption in the CUPS scheduler, which left it vulnerable to attacks against SSLv3. To solve this issue, the **cupsd.conf SSLOptions** keyword has been extended to include two new options, **AllowRC4** and **AllowSSL3**, each of which enables the named feature in **cupsd**. The new options are also supported in the **/etc/cups/client.conf** file. The default is now to disable both RC4 and SSL3 for **cupsd**.

cups now allows underscore in printer names

The **cups** service now allows users to include the underscore character (_) in local printer names.

Unneeded dependency removed from the tftp-server package

Previously, an additional package was installed by default when installing the *tftp-server* package. With this update, the superfluous package dependency has been removed, and the unneeded package is no longer installed by default when installing *tftp-server*.

The deprecated /etc/sysconfig/conman file has been removed

Before introducing the **systemd** manager, various limits for services could be configured in the **/etc/sysconfig/conman** file. After migrating to **systemd**, **/etc/sysconfig/conman** is no longer used and therefore it was removed. To set limits and other daemon parameters, such as LimitCPU=, LimitDATA=, or LimitCORE=, edit the **conman.service** file. For more information, see the systemd.exec(5) manual page. In addition, a new variable LimitNOFILE=10000 has been added to the **systemd.service** file. This variable is commented out by default. Note that after making any changes to the **systemd** configuration, the **systemctl daemon-reload** command must be executed for changes to take effect.

mod_nss rebase to version 1.0.11

The *mod_nss packages* have been upgraded to upstream version 1.0.11, which provides a number of bug fixes and enhancements over the previous version. Notably, **mod_nss** can now enable TLSv1.2, and SSLv2 has been completely removed. Also, support for the ciphers generally considered to be most secure has been added.

The vsftpd daemon now supports DHE and ECDHE cipher suites

The **vsftpd** daemon now supports cipher suites based on the Diffie–Hellman Exchange (DHE) and Elliptic Curve Diffie–Hellman Exchange (ECDHE) key-exchange protocol.

Permissions can now be set for files uploaded with sftp

Inconsistent user environments and strict **umask** settings could result in inaccessible files when uploading using the **sftp** utility. With this update, the administrator is able to force exact permissions for files uploaded using **sftp**, thus avoiding the described issue.

LDAP queries used by ssh-ldap-helper can now be adjusted

Not all LDAP servers use a default schema as expected by the **ssh-ldap-helper** tool. This update makes it possible for the administrator to adjust the LDAP query used by **ssh-ldap-helper** to get public keys from servers using a different schema. Default functionality stays untouched.

A new `createolddir` directive in the logrotate utility

A new logrotate **createolddir** directive has been added to enable automatic creation of the **olddir** directory. For more information, see the logrotate(8) manual page.

Error messages from `/etc/cron.daily/logrotate` are no longer redirected to `/dev/null`

Error messages generated by the daily cronjob of **logrotate** are now sent to the **root** user instead of being silently discarded. In addition, the **/etc/cron.daily/logrotate** script is marked as a configuration file in RPM.

SEED and IDEA based algorithms restricted in `mod_ssl`

The set of cipher suites enabled by default in the **mod_ssl** module of the Apache HTTP Server has been restricted to improve security. SEED and IDEA based encryption algorithms are no longer enabled in the default configuration of **mod_ssl**.

Apache HTTP Server now supports UPN

Names stored in the **subject alternative name** portion of SSL/TLS client certificates, such as the Microsoft User Principle Name, can now be used from the SSLUserName directive and are now available in **mod_ssl** environment variables. Users can now authenticate with their Common Access Card (CAC) or certificate with a UPN in it, and have their UPN used as authenticated user information, consumed by both the access control in Apache and using the **REMOTE_USER** environment variable or a similar mechanism in applications. As a result, users can now set **SSLUserName SSL_CLIENT_SAN_OTHER_msUPN_0** for authentication using UPN.

The `mod_dav` lock database is now enabled by default in the `mod_dav_fs` module

The **mod_dav** lock database is now enabled by default if the Apache HTTP **mod_dav_fs** module is loaded. The default location **ServerRoot/davlockdb** can be overridden using the **DAVLockDB** configuration directive.

`mod_proxy_wstunnel` now supports WebSockets

The Apache HTTP **mod_proxy_wstunnel** module is now enabled by default and it includes support for SSL connections in the **wss://** scheme. Additionally, it is possible to use the **ws://** scheme in the **mod_rewrite** directives. This allows for using WebSockets as a target to **mod_rewrite** and enabling WebSockets in the proxy module.

Chapter 14. Storage

DM rebase to version 4.2

Device Mapper (DM) has been upgraded to upstream version 4.2, which provides a number of bug fixes and enhancements over the previous version including a significant DM crypt performance update and DM core update to support Multi-Queue Block I/O Queueing Mechanism (blk-mq).

Multiqueue I/O scheduling with blk-mq

Red Hat Enterprise Linux 7.2 includes a new multiple queue I/O scheduling mechanism for block devices known as blk-mq. It can improve performance by allowing certain device drivers to map I/O requests to multiple hardware or software queues. The improved performance comes from reducing lock contention present when multiple threads of execution perform I/O to a single device. Newer devices, such as Non-Volatile Memory Express (NVMe), are best positioned to take advantage of this feature due to their native support for multiple hardware submission and completion queues, and their low-latency performance characteristics. Performance gains, as always, will depend on the exact hardware and workload.

The blk-mq feature is currently implemented, and enabled by default, in the following drivers: virtio-blk, mtip32xx, nvme, and rbd.

The related feature, scsi-mq, allows Small Computer System Interface (SCSI) device drivers to use the blk-mq infrastructure. The scsi-mq feature is provided as a Technology Preview in Red Hat Enterprise Linux 7.2. To enable scsi-mq, specify **`scsi_mod.use_blk_mq=y`** on the kernel command line. The default value is **`n`** (disabled).

The device mapper (DM) multipath target, which uses request-based DM, can also be configured to use the blk-mq infrastructure if the **`dm_mod.use_blk_mq=y`** kernel option is specified. The default value is **`n`** (disabled).

It may be beneficial to set **`dm_mod.use_blk_mq=y`** if the underlying SCSI devices are also using blk-mq, as doing so reduces locking overhead at the DM layer.

To determine whether DM multipath is using blk-mq on a system, cat the file **`/sys/block/dm-X/dm/use_blk_mq`**, where **`dm-X`** is replaced by the DM multipath device of interest. This file is read-only and reflects what the global value in **`/sys/module/dm_mod/parameters/use_blk_mq`** was at the time the request-based DM multipath device was created.

New delay_watch_checks and delay_wait_checks options in the multipath.conf file

Should a path be unreliable, as when the connection drops in and out frequently, multipathd will still continuously attempt to use that path. The timeout before multipathd realizes that the path is no longer accessible is 300 seconds, which can give the appearance that multipathd has stalled.

To fix this, two new configuration options have been added: delay_watch_checks and delay_wait_checks. Set the delay_watch_checks to how many cycles multipathd is to watch the path for after it comes online. Should the path fail in under that assigned value, multipathd will not use it. multipathd will then rely on the delay_wait_checks option to tell it how many consecutive cycles it must pass until the path becomes valid again. This prevents unreliable paths from immediately being used as soon as they come back online.

New config_dir option in the multipath.conf file

Users were unable to split their configuration between /etc/multipath.conf and other configuration files. This prevented users from setting up one main configuration file for all their machines and keep machine-specific configuration information in separate configuration files for each machine.

To address this, a new config_dir option was added in the multipath.config file. Users must change the config_dir option to either an empty string or a fully qualified directory path name. When set to anything other than an empty string, multipath will read all .conf files in alphabetical order. It will then apply the configurations exactly as if they had been added to the /etc/multipath.conf. If this change is not made, config_dir defaults to /etc/multipath/conf.d.

New dmstats command to display and manage I/O statistics for regions of devices that use the device-mapper driver

The **dmstats** command provides userspace support for device-mapper I/O statistics. This allows a user to create, manage and report I/O counters, metrics and latency histogram data for user-defined arbitrary regions of device-mapper devices. Statistics fields are now available in **dmsetup** reports and the **dmstats** command adds new specialized reporting modes designed for use with statistics information. For information on the **dmstats** command, see the dmstats(8) man page.

LVM Cache

LVM cache has been fully supported since Red Hat Enterprise Linux 7.1. This feature allows users to create logical volumes (LVs) with a small fast device performing as a cache to larger slower devices. Refer to the lvmcache(7) manual page for information on creating cache logical volumes.

Note the following restrictions on the use of cache LVs:

* The cache LV must be a top-level device. It cannot be used as a thin-pool LV, an image of a RAID LV, or any other sub-LV type.

* The cache LV sub-LVs (the origin LV, metadata LV, and data LV) can only be of linear, stripe, or RAID type.

* The properties of the cache LV cannot be changed after creation. To change cache properties, remove the cache as described in lvmcache(7) and recreate it with the desired properties.

New LVM/DM cache policy

A new **smq** dm-cache policy has been written that the reduces memory consumption and improves performance for most use cases. It is now the default cache policy for new LVM cache logical volumes. Users who prefer to use the legacy **mq** cache policy can still do so by supplying the **–cachepolicy** argument when creating the cache logical volume.

LVM systemID

LVM volume groups can now be assigned an owner. The volume group owner is the system ID of a host. Only the host with the given system ID can use the VG. This can benefit volume groups that exist on shared devices, visible to multiple hosts, which are otherwise not protected from concurrent use from multiple hosts. LVM volume groups on shared devices with an assigned system ID are owned by one host and protected from other hosts.

New `lvmpolld` daemon

The **lvmpolld** daemon provides a polling method for long-running LVM commands. When enabled, control of long-running LVM commands is transferred from the original LVM command to the **lvmpolld** daemon. This allows the operation to continue independent of the original LVM command. The **lvmpolld** daemon is enabled by default.

Before the introduction of the **lvmpolld** daemon, any background polling process originating in an lvm2 command initiated inside a **cgroup** of a systemd service could get killed if the main process (the main service) exited in the **cgroup**. This could lead to premature termination of the lvm2 polling process. Additionally, **lvmpolld** helps to prevent spawning lvm2 polling processes querying for progress on the same task multiple times because it tracks the progress for all polling tasks in progress.

For further information on the **lvmpolld** daemon, see the **lvm.conf** configuration file.

Enhancements to LVM selection criteria

The Red Hat Enterprise Linux 7.2 release supports several enhancements to LVM selection criteria. Previously, it was possible to use selection criteria only for reporting commands; LVM now supports selection criteria for several LVM processing commands as well. Additionally, there are several changes in this release to provide better support for time reporting fields and selection.

For information on the implementation of these new features, see the **LVM Selection Criteria** appendix in the Logical Volume Administration manual.

The default maximum number of SCSI LUNs is increased

The default value for the **max_report_luns** parameter has been increased from 511 to 16393. This parameter specifies the maximum number of logical units that may be configured when the systems scans the SCSI interconnect using the Report LUNs mechanism.

Chapter 15. System and Subscription Management

PowerTOP now respects user-defined report file names

Previously, PowerTOP report file names were generated in an unclear, undocumented way. With this update, the implementation has been improved, and the generated file names now respect the names requested by the user. This applies to both CSV and HTML reports.

Amended `yum-config-manager` commands

Previously, running the **`yum-config-manager --disable`** command disabled all configured repositories, while the **`yum-config-manager --enable`** command did not enable any. This inconsistency has been fixed. The **`--disable`** and **`--enable`** commands now require the use of '*' in the syntax, and **`yum-config-manager --enable \*`** enables repositories. Running the commands without the addition of '*' prints a message asking the user to run **`yum-config-manager --disable \*`** or **`yum-config-manager --enable \*`** if they want to disable or enable repositories.

New search-disabled-repos plug-in for yum

The search-disabled-repos plug-in for yum has been added to the subscription-manager packages. This plug-in allows users to successfully complete yum operations that fail due to the source repository being dependent on a disabled repository. When search-disabled-repos is installed in the described scenario, yum displays instructions to temporarily enable repositories that are currently disabled and to search for missing dependencies.

If you choose to follow the instructions and turn off the default notify_only behavior in the /etc/yum/pluginconf.d/search-disabled-repos.conf file, future yum operations will prompt you to temporarily or permanently enable all the disabled repositories needed to fulfill the yum transaction.

Acquiring hypervisor data in parallel

With this update, virt-who is able to acquire data from multiple hypervisors in parallel. Previously, virt-who could read data only from a single hypervisor at a time, and if one hypervisor in a series was nonfunctional, virt-who waited for its response and thus failed. Reading parallel hypervisors works around this problem and prevents the described failure.

Filtering for hypervisors reported by virt-who

The virt-who service introduces a filtering mechanism for the Subscription Manager reports. As a result, users can now choose which hosts virt-who should display according to the specified parameters. For example, they can filter out hosts that do not run any Red Hat Enterprise Linux guests, or hosts that run guests of a specified version of Red Hat Enterprise Linux.

Improved visualization of host-to-guest association

The **`-p`** option has been added to the virt-who utility. When used with **`-p`**, virt-who output displays Javascript Object Notation (JSON)-encoded map of the host-guest association. In addition, the information on host-guest association logged in the /var/log/rhsm/rhsm.log file is now formatted in JSON as well.

virt-who output displayed as host names

It is now possible to configure the virt-who query so that its results are displayed as host names instead of as Universally Unique Identifiers (UUIDs) when viewed in Red Hat Satellite and Red Hat Customer Portal. To enable the function, add the **`hypervisor_id=hostname`** option to the configuration file in the /etc/virt-who.d/ directory. Ideally, this should be done before using virt-who for the first time, otherwise changing the configuration duplicates the hypervisor.

Pre-filled virt-who configuration file

A default configuration file for virt-who has been placed in the /etc/virt-who.d/ directory. It contains a template and instructions for the user to configure virt-who. This replaces the deprecated configuration that uses the /etc/sysconfig/virt-who file.

Enhanced proxy connection options

With Red Hat Enterprise Linux 7.2, the virt-who utility can handle the HTTP_PROXY and HTTPS_PROXY environment variables, and thus correctly uses the proxy server when requested. This allows virt-who to connect to the Hyper-V hypervisor and Red Hat Enterprise Virtualization Manager through proxy.

Subscription Manager now supports syslog

The *subscription-manager* tool can now use the syslog as the log handler and formatter in addition to separate log used previously. The handler and formatter is configured in the **`/etc/rhsm/logging.conf`** configuration file.

Subscription Manager is now part of Initial Setup

The Subscription Manager component of Firstboot has been ported to the Initial Setup utility. Users are now able to register the system from the main menu of Initial Setup after installing a Red Hat Enterprise Linux 7 system and rebooting for the first time.

Subscription Manager now displays the server URL when registering on a command line

When registering a system using the **`subscription-manager`** command on a command line, the tool now also shows the URL of the server being used for the registration when asking for user name and password. This helps the user determine which credentials to use.

Manage Repositories dialog in Subscription Manager is now more responsive

The Manage Repositories dialog in the graphical version of Subscription Manager (the *subscription-manager-gui* package) has been updated to no longer fetch information on each checkbox change. Instead, the system state is only synchronized when the new **`save`** button is clicked. This removes delays users experienced in previous versions caused by the system state being updated on each checkbox action, and repository management is now significantly more responsive.

Chapter 16. Virtualization

qemu-kvm supports virtual machine shutdown trace events

Support has been added for qemu-kvm trace events during the virtual machine system shutdown process, which allows users to get detailed diagnostics about a guest system's shutdown requests issued by the **`virsh shutdown`** command or by the virt-manager application. This provides users with enhanced capabilities for isolating and debugging KVM guest problems during shutdown.

Intel MPX exposed to the guest

With this update, qemu-kvm allows the Intel Memory Protection Extensions (MPX) feature to be exposed to the guest. On the Intel 64 host systems that support MPX, this enables the use of a set of extensions that provide hardware support for bounds protection on pointer references.

Guest memory dump extraction from the qemu-kvm core

The dump-guest-memory.py script has been introduced into QEMU, which makes it possible to analyze a guest memory dump from the qemu-kvm core in case of a guest kernel failure. For further information, see the related help text by using the **`help dump-guest-memory`** command.

virt-v2v is fully supported

With Red Hat Enterprise Linux 7.2, the virt-v2v command-line tool has become fully supported. This tool converts virtual machines running on foreign hypervisors to run on KVM. Currently, virt-v2v can convert Red Hat Enterprise Linux and Windows guests running on Red Hat Enterprise Linux 5 Xen and VMware vCenter.

Virtualization on IBM Power Systems

Red Hat Enterprise Linux with KVM is supported on AMD64 and Intel 64 systems, but not on IBM Power Systems. Red Hat currently provides a POWER8-based solution with Red Hat Enterprise Virtualization for IBM Power Systems.

More information on version support and installation procedures can be found in the following Knowledge Base article: https://access.redhat.com/articles/1247773.

Hyper-V TRIM support

Now it is possible to use Thin Provisioned Hyper-V virtual hard disk (VHDX). The update adds support to shrink the underlining VHDX files for Microsoft Hyper-V virtual machines to actual used size.

KVM support for tcmalloc

KVM can now use the tcmalloc library, which provides a significant performance improvement in I/O operations per second.

Chapter 17. Atomic Host and Containers

Red Hat Enterprise Linux Atomic Host 7.2

Red Hat Enterprise Linux Atomic Host is a secure, lightweight, and minimal-footprint operating system optimized to run Linux containers.

It is pre-installed with the following tools to support Linux containers:

* Docker - an open source engine that automates the deployment of any application as a lightweight, portable, self-sufficient container that will run virtually anywhere

* atomic - defines the entrypoint for Atomic hosts

* kubernetes - provides container cluster management

* etcd - provides a highly-available key value store for shared configuration

* flannel - contains an etcd-driven address management agent, which manages IP addresses of overlay networks between systems running containers that need to communicate with one another

Red Hat Enterprise Linux Atomic Host makes use of the following technologies:

* OSTree and rpm-OSTree - These projects provide atomic upgrades and rollback capability

* systemd - a new init system for Linux that enables faster boot times and easier orchestration

* SELinux - enabled by default to provide complete multi-tenant security

Also, **`Cockpit`** is available on Red Hat Enterprise Linux as a separate Extras package and on Red Hat Enterprise Linux Atomic Host, as a Container Image, **`cockpit-ws`**. Cockpit is a server administration interface that makes it easy to administer Red Hat Enterprise Linux servers through a web browser.

Notable changes in Red Hat Enterprise Linux Atomic Host 7.2 include:

OSTree update: For the full list of updated packages, see https://access.redhat.com/articles/2050783.

Updated packages:

* docker-1.8.2-8.el7

* flannel-0.5.3-8.el7

* cockpit-0.77-3.1.el7

* storaged-2.2.0-3.el7

* kubernetes-1.0.3-0.2.gitb9a88a7.el7

* atomic-1.6-6.gitca1e384.el7

* python-websocket-client-0.32.0-116.el7

* python-docker-py-1.4.0-118.el7

New packages:

* docker-distribution-2.1.1-3.el7

Updated container images:

Red Hat Enterprise Linux 7.2 Container Image Update

Red Hat Enterprise Linux Atomic rsyslog Container Image

Red Hat Enterprise Linux Atomic sadc Container Image

Red Hat Enterprise Linux Atomic Tools Container Image

Red Hat Enterprise Linux Atomic cockpit-ws Container Image

New container images:

Red Hat Enterprise Linux Etcd Container Image

Red Hat Enterprise Linux Kubernetes-controller Container Image

Red Hat Enterprise Linux Kubernetes-apiserver Container Image

Red Hat Enterprise Linux Kubernetes-scheduler Container Imafe

Updates to container-related packages

docker-1.8.2-8.el7

The *docker* packages have been upgraded to upstream version 1.8.2.

Additionally, *docker* also includes the following changes:

* **`docker`** now displays a warning message if you are using the loopback device as a backend storage option.

* The **`docker info`** command now shows the rpm version of the client and server.

* The default mount propagation is **`Slave`** instead of **`Private`**. This allows volume (bind) mounts, to be altered on the host and the new mounts show up inside of the container.

* The **`--add-registry`** and **`--block-registry`** options have been added. This allows additional registries to be specified in addition to **`docker.io`**.

* You can now inspect the content of remote repositories and check for newer versions. This functionality is implemented in the **`atomic verify`** command.

* Improvements to system logging, now all access to the docker daemon is logged.

flannel-0.5.3-8.el7

* flannel's network prefix was changed from **`coreos.com/network`** to **`atomic.io/network`**.

* flannel's behaviour when the first ping packet was lost has been fixed.

* The **`flanneld.service`** now starts after the network is ready.

kubernetes-1.0.3-0.2.gitb9a88a7.el7

* **`kubectl version`** now displays the correct version.

* When running **`kube-apiserver`** on port 443 in secure mode, some capabilities are missing. As a workaround, the **`kube-apiserver`** binary has to be modified by running:

```
# chown root:root /usr/bin/kube-apiserver
```

```
# chmod 700 /usr/bin/kube-apiserver

# setcap CAP_NET_BIND_SERVICE=ep /usr/bin/kube-apiserver
```

cockpit-0.77-3.1.el7

* Cockpit now displays the limit for the number of supported hosts when adding servers to the dashboard.

* Cleaner bookmarkable URLs.

* Includes basic SSH key authentication functionality.

* Basic interactions with multipath storage have been fixed.

* When password authorization is not possible, Cockpit displays an informative message.

* Authentication now works when embedding Cockpit.

Removed systemd socket activation

For security reasons, systemd socket activation, which was supported in earlier versions of Docker, has been removed. Now, it is not recommended to use the docker group as a mechanism for talking to the docker daemon as a non-privileged user. Instead, set up sudo for this type of access. If the docker daemon is not running after the upgrade, create the **`/etc/sysconfig/docker.rpmnew`** file, add any local customization to it and replace **`/etc/sysconfig/docker`** with it. Additionally, remove the **`-H fd://`** line from **`/etc/sysconfig/docker`** if it is present.

Chapter 18. Red Hat Software Collections

Red Hat Software Collections is a Red Hat content set that provides a set of dynamic programming languages, database servers, and related packages that you can install and use on all supported releases of Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 on AMD64 and Intel 64 architectures.

Dynamic languages, database servers, and other tools distributed with Red Hat Software Collections do not replace the default system tools provided with Red Hat Enterprise Linux, nor are they used in preference to these tools. Red Hat Software Collections uses an alternative packaging mechanism based on the **scl** utility to provide a parallel set of packages. This set allows for optional use of alternative package versions on Red Hat Enterprise Linux. By using the **scl** utility, users can pick and choose which package version they want to run at any time.

Red Hat Developer Toolset is now a part of Red Hat Software Collections. It is included as a separate Software Collection. Red Hat Developer Toolset is designed for developers working on the Red Hat Enterprise Linux platform. It provides current versions of the GNU Compiler Collection, GNU Debugger, Eclipse development platform, and other development, debugging, and performance monitoring tools.

Important

Red Hat Software Collections has a shorter life cycle and support term than Red Hat Enterprise Linux. For more information, see the Red Hat Software Collections Product Life Cycle.

See the Red Hat Software Collections documentation for the components included in the set, system requirements, known problems, usage, and specifics of individual Software Collections.

See the Red Hat Developer Toolset documentation for more information about the components included in this Software Collection, installation, usage, known problems, and more.

Part II. Technology Previews

This part provides an overview of Technology Previews introduced or updated in Red Hat Enterprise Linux 7.2.

For more information on Red Hat Technology Previews, see https://access.redhat.com/support/offerings/techpreview/.

Chapter 19. Authentication and Interoperability

Use of AD and LDAP sudo providers

The Active Directory (AD) provider is a back end used to connect to an AD server. In Red Hat Enterprise Linux 7.2, using the AD sudo provider together with the LDAP provider is supported as a Technology Preview. To enable the AD sudo provider, add the **`sudo_provider=ad`** setting in the [domain] section of the **`sssd.conf`** file.

DNSSEC available as Technology Preview in Identity Management

Identity Management servers with integrated DNS now support DNS Security Extensions (DNSSEC), a set of extensions to DNS that enhance security of the DNS protocol. DNS zones hosted on Identity Management 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: http://tools.ietf.org/html/rfc6781#section-2

Secure Domain Name System (DNS) Deployment Guide: http://dx.doi.org/10.6028/NIST.SP.800-81-2

DNSSEC Key Rollover Timing Considerations: http://tools.ietf.org/html/rfc7583

Note that Identity Management 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 described in the Red Hat Enterprise Linux Networking Guide: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Networking_Guide/ch-Configure_Host_Names.html#sec-Recommended_Naming_Practices.

Nunc Stans event framework available for Directory Server

A new Nunc Stans event framework to handle multiple simultaneous connections has been added as Technology Preview. The framework allows supporting several thousand active connections with no performance degradation. It is disabled by default.

Browser for the JSON-RPC API in IdM is available

This update implements a browser for the JSON-RPC API in Identity Management. The browser can be used to view the API. Note that this feature is experimental and the API is not yet supported.

New packages: ipsilon

The *ipsilon* packages provide the Ipsilon identity provider service for federated single sign-on (SSO). Ipsilon links authentication providers and applications or utilities to allow for SSO. It includes a server and utilities to configure Apache-based service providers.

The Ipsilon server and toolkit is designed to configure Apache-based identity Service Providers. The server is a pluggable self-contained **`mod_wsgi`** application that provides federated SSO to web applications.

Ipsilon is introduced in this release as a Technology Preview. Customers are advised not to consider integration of this service for production environments at this time.

Chapter 20. File Systems

OverlayFS

OverlayFS is a type of union file system. It allows the user 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. Refer to the kernel file Documentation/filesystems/overlayfs.txt for additional information.

OverlayFS remains a Technology Preview in Red Hat Enterprise Linux 7.2 under most circumstances. As such, the kernel will log warnings when this technology is activated.

Full support is available for OverlayFS when used with Docker under the following restrictions:

* OverlayFS is only supported for use as a Docker graph driver. Its use can only be supported for container COW content, not for persistent storage. Any persistent storage must be placed on non-OverlayFS volumes to be supported. Only default Docker configuration can be used; that is, 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.

* SELinux must be enabled and in enforcing mode on the physical machine, but must be disabled in the container when performing container separation; that is, /etc/sysconfig/docker must not contain --selinux-enabled. SELinux support for OverlayFS is being worked on upstream, and is expected in a future release.

* The OverlayFS kernel ABI and userspace behavior are not considered stable, and may see changes in future updates.

* In order to make the yum and rpm utilities work properly inside the container, the user should be using the yum-plugin-ovl packages.

Note that OverlayFS provides a restricted set of the POSIX standards. Test your application thoroughly before deploying it with OverlayFS.

There are also several known issues associated with OverlayFS as of Red Hat Enterprise Linux 7.2 release. For details, see 'Non-standard behavior' in the Documentation/filesystems/overlayfs.txt file.

Support for NFSv4 clients with flexible file layout

Red Hat Enterprise Linux 7.2 adds support for flexible file layout on NFSv4 clients. This technology enables advanced features such as non-disruptive file mobility and client-side mirroring, providing enhanced usability in areas such as databases, big data and virtualization.

See https://datatracker.ietf.org/doc/draft-ietf-nfsv4-flex-files/ for detailed information about NFS flexible file layout.

Btrfs file system

The Btrfs (B-Tree) file system is supported as a Technology Preview in Red Hat Enterprise Linux 7.2. This file system offers advanced management, reliability, and scalability features. It enables users to create snapshots, it enables compression and integrated device management.

pNFS Block Layout Support

As a Technology Preview, the upstream code has been backported to the Red Hat Enterprise Linux client to provide pNFS block layout support.

Chapter 21. Hardware Enablement

Runtime Instrumentation for IBM System z

Support for the Runtime Instrumentation feature is available as a Technology Preview in Red Hat Enterprise Linux 7.2 on IBM System z. Runtime Instrumentation enables advanced analysis and execution for a number of user-space applications available with the IBM zEnterprise EC12 system.

LSI Syncro CS HA-DAS adapters

Red Hat Enterprise Linux 7.1 included code in the megaraid_sas driver to enable LSI Syncro CS high-availability direct-attached storage (HA-DAS) adapters. While the megaraid_sas driver is fully supported for previously enabled adapters, the use of this driver for Syncro CS is available as a Technology Preview. Support for this adapter is provided directly by LSI, your system integrator, or system vendor. Users deploying Syncro CS on Red Hat Enterprise Linux 7.2 are encouraged to provide feedback to Red Hat and LSI. For more information on LSI Syncro CS solutions, please visit http://www.lsi.com/products/shared-das/pages/default.aspx.

Chapter 22. Kernel

Multiple CPU support in kdump on AMD64 and Intel 64 systems

On AMD64 and Intel 64 systems, the **kdump** kernel crash dumping mechanism can now boot with more than one CPU enabled. This solves a problem on systems with large memory size where, due to high input and output when creating a kernel crash dump, Linux could fail to allocate interrupts for devices when only one CPU was enabled: **maxcpus=1** or **nr_cpus=1**.

To enable multiple CPUs in the crash kernel, provide the **nr_cpus=X** (where **X** is the number of processors) and **disable_cpu_apicid=0** options on the kernel command line.

The criu tool

Red Hat Enterprise Linux 7.2 introduces the **criu** tool as a Technology Preview. This tool implements **Checkpoint/Restore in User-space**, which can be used to freeze a running application and store it as a collection of files. Later, the application can be restored from its frozen state.

The **criu** tool depends on **Protocol Buffers**, a language-neutral, platform-neutral extensible mechanism for serializing structured data. The *protobuf* and *protobuf-c* packages, which provide this dependency, are also added to Red Hat Enterprise Linux 7.2 as a Technology Preview.

User namespace

This feature provides additional security to servers running Linux containers by providing better isolation between the host and the containers. Administrators of a container are no longer able to perform administrative operations on the host, which increases security.

LPAR Watchdog for IBM System z

An enhanced watchdog driver for IBM System z is available as a Technology Preview. This driver supports Linux logical partitions (LPAR) as well as Linux guests in the z/VM hypervisor, and provides automatic reboot and automatic dump capabilities if a Linux system becomes unresponsive.

Dynamic kernel updates with kpatch

The kpatch utility allows users to manage a collection of binary kernel patches which can be used to dynamically patch the kernel without rebooting. kpatch is supported as a Technology Preview and only for AMD64 and Intel 64 architectures.

i40evf handles big resets

The most common type of reset that the Virtual Function (VF) encounters is a Physical Function (PF) reset that cascades down into a VF reset for each VF. However, for 'bigger' resets, such as a Core or EMP reset, when the device is reinitialized, the VF previously did not get the same VSI, so the VF was not able to recover, as it continued to request resources for its original VSI. As a Technology Preview, this update adds an extra state to the admin queue state machine, so that the driver can re-request its configuration information at runtime. During reset recovery, this bit is set in the aq_required field, and the configuration information is fetched before attempting to bring the driver back up.

Support for OPA kernel driver

Intel Omni-Path Architecture (OPA) kernel driver, which is supported as a Technology Preview, provides Host Fabric Interconnect (HFI) hardware with initialization and setup for High Performance Data transfers to other compute and I/O nodes in a High Performance Computing (HPC) cluster.

Support for Diag0c on IBM System z

As a Technology Preview, Red Hat Enterprise Linux 7.2 introduces support for the Diag0c feature on IBM System z. Diag0c support makes it possible to read the CPU performance metrics provided by the z/VM hypervisor, and allows obtaining the management time for each online CPU of a Linux guest where the diagnose task is executed.

Chapter 23. Networking

i40e and i40evf rebase to versions 1.3.21-k and 1.3.13

The i40e and i40evf kernel drivers have been updated to version 1.3.21-k and 1.3.13. These updated drivers are included as a Technology Preview in Red Hat Enterprise Linux 7.2.

On i40e ports, an attempt to run iSCSI related commands previously led to loss of network connectivity out of i40e ports. This update fixes this bug, and the system now allows for iSCSI commands to proceed.

Cisco usNIC driver

Cisco Unified Communication Manager (UCM) servers have an optional feature to provide a Cisco proprietary User Space Network Interface Controller (usNIC), which allows performing Remote Direct Memory Access (RDMA)-like operations for user-space applications. The libusnic_verbs driver, which is supported as a Technology Preview, makes it possible to use usNIC devices via standard InfiniBand RDMA programming based on the Verbs API.

Cisco VIC kernel driver

The Cisco VIC Infiniband kernel driver, which is supported as a Technology Preview, allows the use of Remote Directory Memory Access (RDMA)-like semantics on proprietary Cisco architectures.

Trusted Network Connect

Trusted Network Connect, supported as a Technology Preview, is used with existing network access control (NAC) solutions, such as TLS, 802.1X, or IPsec to integrate endpoint posture assessment; that is, collecting an endpoint's system information (such as operating system configuration settings, installed packages, and others, termed as integrity measurements). Trusted Network Connect is used to verify these measurements against network access policies before allowing the endpoint to access the network.

SR-IOV functionality in the qlcnic driver

Support for Single-Root I/O virtualization (SR-IOV) has been added to the qlcnic driver as a Technology Preview. Support for this functionality will be provided directly by QLogic, and customers are encouraged to provide feedback to QLogic and Red Hat. Other functionality in the qlcnic driver remains fully supported.

Chapter 24. Storage

Multi-queue I/O scheduling for SCSI

Red Hat Enterprise Linux 7.2 includes a new multiple-queue I/O scheduling mechanism for block devices known as blk-mq. The scsi-mq package allows the Small Computer System Interface (SCSI) subsystem to make use of this new queuing mechanism. This functionality is provided as a Technology Preview and is not enabled by default. To enable it, add **`scsi_mod.use_blk_mq=Y`** to the kernel command line.

Improved LVM locking infrastructure

`lvmlockd` is a next generation locking infrastucture for LVM. It allows LVM to safely manage shared storage from multiple hosts, using either the **`dlm`** or **`sanlock`** lock managers. **`sanlock`** allows **`lvmlockd`** to coordinate hosts through storage-based locking, without the need for an entire cluster infrastructure. For more information, see the **`lvmlockd`**(8) man page.

Targetd plug-in from the libStorageMgmt API

Since Red Hat Enterprise Linux 7.1, storage array management with libStorageMgmt, a storage array independent API, has been fully supported. The provided API is stable, consistent, and allows developers to programmatically manage different storage arrays and utilize the hardware-accelerated features provided. System administrators can also use libStorageMgmt to manually configure storage and to automate storage management tasks with the included command-line interface.

The Targetd plug-in is not fully supported and remains a Technology Preview.

DIF/DIX

DIF/DIX is a new addition to the SCSI Standard. It is fully supported in Red Hat Enterprise Linux 7.2 for the HBAs and storage arrays specified in the Features chapter, but it remains in Technology Preview for all other HBAs and storage arrays.

DIF/DIX increases the size of the commonly used 512 byte disk block from 512 to 520 bytes, adding the Data Integrity Field (DIF). The DIF stores a checksum value for the data block that is calculated by the Host Bus Adapter (HBA) when a write occurs. The storage device then confirms the checksum on receipt, and stores both the data and the checksum. Conversely, when a read occurs, the checksum can be verified by the storage device, and by the receiving HBA.

Chapter 25. Virtualization

Nested virtualization

As a Technology Preview, Red Hat Enterprise Linux 7.2 offers nested virtualization. This feature enables KVM to launch guests that can act as hypervisors and create their own guests.

The virt-p2v tool

Red Hat Enterprise Linux 7.2 offers the virt-p2v tool as a Technology Preview. virt-p2v (physical to virtual) is a CD-ROM, ISO or PXE image that the user can boot on a physical machine, and that creates a KVM virtual machine with disk contents identical to the physical machine.

USB 3.0 support for KVM guests

USB 3.0 host adapter (xHCI) emulation for KVM guests remains a Technology Preview in Red Hat Enterprise Linux 7.2.

VirtIO-1 support

Virtio drivers have been updated to Kernel 4.1 to provide VirtIO 1.0 Device Support.

Part III. Device Drivers

This chapter provides a comprehensive listing of all device drivers which were updated in Red Hat Enterprise Linux 7.2.

Chapter 26. Storage Driver Updates

- The hpsa driver has been updated to version 3.4.4-1-RH4.
- The qla2xxx driver has been updated to version 8.07.00.18.07.2-k.
- The lpfc driver has been updated to version 10.7.0.1.
- The megaraid_sas driver has been updated to version 06.807.10.00.
- The fnic driver has been updated to version 1.6.0.17.
- The mpt2sas driver has been updated to version 20.100.00.00.
- The mpt3sas driver has been updated to version 9.100.00.00.
- The Emulex be2iscsi driver has been updated to version 10.6.0.0r.
- The aacraid driver has been updated to version 1.2.
- The bnx2i driver has been updated to version 2.7.10.1.
- The bnx2fc driver has been updated to version 2.4.2.

Chapter 27. Network Driver Updates

- The tg3 driver has been updated to version 3.137.
- The e1000 driver has been updated to version 7.3.21-k8-NAPI, which provides support for txtd update delay when using the xmit_more Boolean variable.
- The e1000e driver has been updated to version 3.2.5-k.
- The igb driver has been updated to version 5.2.15-k.
- The igbvf driver has been updated to version 2.0.2-k.
- The ixgbevf driver has been updated to version 2.12.1-k.
- The ixgbe driver has been updated to version 4.0.1-k.
- The bna driver and firmware have been updated to version 3.2.23.0r.
- The bnx2 driver has been updated to version 2.2.6.
- The CNIC driver has been updated to version 2.5.21.
- The bnx2x driver has been updated to version 1.710.51-0, which also adds qlogic NPAR support for qlogic-nx2 adapters.
- The be2net driver has been updated to version 10.6.0.3r.
- The qlcnic driver has been updated to version 5.3.62.
- The qlge driver has been updated to version 1.00.00.34. It fixes a race condition between the New API (NAPI) registration and unregistration which previously led to the system crash. This race condition occured if certain parameters were changed while the Network Interface Card (NIC) was set to "down".
- The r8169 driver has been updated to version 2.3LK-NAPI.
- The i40e driver has been updated to version 1.3.21-k.
- The i40evf driver has been updated to version 1.3.13.
- The netxen_nic driver has been updated to version 4.0.82.
- The sfc driver has been updated to the latest upstream version.
- This update adds the fm10k driver of version 0.15.2-k.
- This update adds VTI6 support including netns capabilities.
- The bonding driver has been updated to version 3.7.1.
- The iwlwifi driver has been updated to the latest upstream version.
- The vxlan driver has been updated to version 0.1.

Chapter 28. Graphics Driver and Miscellaneous Driver Updates

- The HDA driver has been updated to the latest upstream version to use the new jack kctls method.
- The HPI driver has been updated to version 4.14.
- The Realtek HD-audio codec driver has been updated to include the update of EAPD init codes.
- The IPMI driver has been updated to replace the timespec usage by timespec64.
- The i915 driver has been updated to include the rebase of ACPI Video Extensions driver in Red Hat Enterprise Linux 7.2.
- The ACPI Fan driver has been updated to version 0.25.
- The Update NVM-Express driver has been updated to version 3.19.
- The rtsx driver has been updated to version 4.0 to support rtl8402, rts524A, rts525A chips.
- The Generic WorkQueue Engine device driver has been updated to the latest upstream version.
- The PCI driver has been updated to version 3.16.
- The EDAC kernel module has been updated to provide support for Intel Xeon v4 processors.
- The pstate driver has been updated to support 6th Generation Intel Core processor.
- The intel_idle driver has been updated to support 6th Generation Intel Core processor.

Part IV. Deprecated Functionality

This part provides an overview of functionality that has been deprecated as of Red Hat Enterprise Linux 7.2 or is planned to be deprecated in the future.

Chapter 29. Deprecated Functionality as of Red Hat Enterprise Linux 7.2

No functionality of Red Hat Enterprise Linux 7 has been deprecated in this release.

Chapter 30. Deprecated Functionality in Future Releases

The following Emulex boards will be deprecated *after* Red Hat Enterprise Linux 7 and will likely not be supported in the next major release:

Table 30.1. Deprecated Emulex Boards

BladeEngine 2 (BE2) Devices	
0x0704	TIGERSHARK
Fibre Channel (FC) Devices	
0x1ae5	FIREFLY
0xe100	PROTEUS_VF
0xe131	BALIUS
0xe180	PROTEUS_PF
0xf095	RFLY
0xf098	PFLY
0xf0a1	LP101
0xf0a5	TFLY
0xf0d1	BSMB
0xf0d5	BMID
0xf0e1	ZSMB
0xf0e5	ZMID
0xf0f5	NEPTUNE
0xf0f6	NEPTUNE_SCSP
0xf0f7	NEPTUNE_DCSP
0xf180	FALCON
0xf700	SUPERFLY
0xf800	DRAGONFLY
0xf900	CENTAUR
0xf980	PEGASUS
0xfa00	THOR
0xfb00	VIPER
0xfc00	LP10000S
0xfc10	LP11000S
0xfc20	LPE11000S
0xfc50	PROTEUS_S
0xfd00	HELIOS
0xfd11	HELIOS_SCSP
0xfd12	HELIOS_DCSP
0xfe00	ZEPHYR
0xfe05	HORNET
0xfe11	ZEPHYR_SCSP
0xfe12	ZEPHYR_DCSP

To check the PCI IDs of the hardware on your system, run the **`lspci -nn`** command.

Part V. Known Issues

This part documents known problems in Red Hat Enterprise Linux 7.2.

Chapter 31. General Updates

Upgrading from Red Hat Enterprise Linux 6 may fail on IBM Power Systems

Because of a bug in the **yaboot** boot loader, upgrading from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7 can fail on IBM Power Systems servers with an **Unknown or corrupt filesystem** error.

This problem is typically caused by a misplaced **yaboot.conf** configuration file. Make sure that this file exists, that it is valid, and that it is placed on a standard (non-LVM) /boot partition.

The /etc/os-release file contains outdated information after system upgrade

Upgrading to the next minor release (for example, from Red Hat Enterprise Linux 7.1 to 7.2) does not update the **/etc/os-release** file with the new product number. Instead, this file continues to list the previous release number, and a new file named **os-release.rpmnew** is placed in the **/etc** directory.

If you require the **/etc/os-release** file to be up-to-date, replace it with **/etc/os-release.rpmnew**.

Chapter 32. Authentication and Interoperability

Kerberos ticket requests are refused for short lifetimes

Due to a bug in Active Directory, Kerberos ticket requests for short (generally below three minutes) lifetimes, are refused. To work around this problem, request longer-lived (above five minutes) tickets instead.

Replication from a Red Hat Enterprise Linux 7 machine to a Red Hat Enterprise Linux 6 machine fails

Currently, the Camellia Kerberos encryption types (enctypes) are included as possible default enctypes in the krb5, krb5-libs, krb5-server packages. As a consequence, replication from a Red Hat Enterprise Linux 7 machine to a Red Hat Enterprise Linux 6 machine fails with an error message. To work around this problem, use the default enctype controls, or tell kadmin or ipa-getkeytab which encryption types to use.

A harmless error message is logged on SSSD startup

If SSSD is connected to an IdM server that does not have a trust relationship established with an AD server, the following harmless error message is printed to the SSSD domain log on startup:

Internal Error (Memory buffer error)

To prevent the harmless error message from occurring, set **`subdomains_provider`** to **`none`** in the sssd.conf file if the environment does not expect setting any trusted domains.

DNS zones with recently generated DNSSEC keys are not signed properly

IdM does not properly sign DNS zones with recently generated DNS Security Extensions (DNSSEC) keys. The named-pkcs11 service logs the following error in this situation:

The attribute does not exist: 0x00000002

The bug is caused by a race condition error in the DNSSEC key generation and distribution process. The race condition prevents named-pkcs11 from accessing new DNSSEC keys.

To work around this problem, restart named-pkcs11 on the affected server. After the restart, the DNS zone is properly signed. Note that the bug might reappear after the DNSSEC keys are changed again.

The old realmd version is started when updating realmd while it is running

The **`realmd`** daemon starts only when requested, then performs a given action, and after some time it times out. When **`realmd`** is updated while it is still running, the old version of **`realmd`** starts upon a next request because **`realmd`** is not restarted after the update. To work around this problem, make sure that **`reamld`** is not running before updating it.

ipa-server-install and ipa-replica-install do not validate their options

The **`ipa-server-install`** and **`ipa-replica-install`** utilities do currently not validate the options supplied to them. If the user passes incorrect values to the utilities, the installation fails. To work around the problem, make sure to supply correct values, and then run the utilities again.

Chapter 33. Compiler and Tools

Multiple bugs when booting from SAN over FCoE

Multiple bugs have arisen from the current implementation of boot from Storage Area Network (SAN) using Fibre Channel over Ethernet (FCoE). Red Hat is targeting a future release of Red Hat Enterprise Linux 7 for the fixes for these bugs. For a list of the affected bugs and workarounds (where available), please contact your Red Hat support representative.

Valgrind cannot run programs built against an earlier version of Open MPI

Red Hat Enterprise Linux 7.2 supports only the Open MPI application binary interface (ABI) in version 1.10, which is incompatible with the previously shipped 1.6 version of the Open MPI ABI. As a consequence, programs that are built against the earlier version of Open MPI cannot be run under Valgrind included in Red Hat Enterprise Linux 7.2. To work around this problem, use the Red Hat Developer Toolset version of Valgrind for programs linked against Open MPI version 1.6.

Synthetic functions generated by GCC confuse SystemTap

A GCC optimization can generate synthetic functions for partially inlined copies of other functions. These synthetic functions look like first-class functions and confuse tools such as SystemTap and GDB because SystemTap probes can be placed on both synthetic and real function entry points. This can result in multiple SystemTap probe hits per a single underlying function call.

To work around this problem, a SystemTap script may need to adopt countermeasures, such as detecting recursion and suppressing probes related to inlined partial functions. For example, the following script:

probe kernel.function("can_nice").call { }

could attempt to avoid the described problem as follows:

global in_can_nice% probe kernel.function("can_nice").call { in_can_nice[tid()] ++; if (in_can_nice[tid()] > 1) { next } /* real probe handler here */ } probe kernel.function("can_nice").return { in_can_nice[tid()] --; }

Note that this script does not take into account all possible scenarios. It would not work as expected in case of, for example, missed kprobes or kretprobes, or genuine intended recursion.

SELinux AVC generated when ABRT collects backtraces

If the new, optional ABRT feature that allows collecting backtraces from crashed processes without the need to write a core-dump file to disk is enabled (using the **`CreateCoreBacktrace`** option in the **`/etc/abrt/plugins/CCpp.conf`** configuration file), an SELinux AVC message is generated when the **`abrt-hook-ccpp`** tool tries to use the **`sigchld`** access on a crashing process in order to get the list of functions on the process' stack.

GDB keeps watchpoints active even after reporting them as hit

In some cases, GDB can incorrectly keep watchpoints active even after reporting them as hit. This results in the watchpoints getting hit for the second time, only this time the hardware indication is no longer recognized as a watchpoint and is printed as a generic SIGTRAP signal instead. There are several ways to work around this problem and stop the excessive SIGTRAP reporting.

* Type **`continue`** when seeing a SIGTRAP after a watchpoint has been hit.

* Instruct GDB to ignore the SIGTRAP signal by adding the following line to your **~/.gdbinit** configuration file:

handle SIGTRAP nostop noprint

* Use software watchpoints instead of their hardware equivalents. Note that the debugging is significantly slower with software watchpoints, and only the **watch** command is available (not **rwatch** or **awatch**). Add the following line to your **~/.gdbinit** configuration file:

set can-use-hw-watchpoints 0

* Use the Red Hat Developer Toolset 4.0 version of GDB.

Booting fails using grubaa64.efi

Due to issues in pxeboot or the PXE configuration file, installing Red Hat Enterprise Linux 7.2 using the 7.2 grubaa64.efi boot loader either fails or experiences significant delay in booting the operating system. As a workaround, use the 7.1 grubaa64.efi file instead of the 7.2 grubaa64.efi file when installing Red Hat Enterprise Linux 7.2.

MPX feature in GCC requires Red Hat Developer Toolset version of the libmpx library

The libmpxwrappers library is missing in the gcc-libraries version of the libmpx library. As a consequence, the Memory Protection Extensions (MPX) feature might not work correctly in GCC, and the application might not link properly. To work around this problem, use the Red Hat Developer Toolset 4.0 version of the libmpx library.

Chapter 34. Desktop

Broken pygobject3 package dependencies prevent upgrade from Red Hat Enterprise Linux 7.1

The *pygobject3-devel.i686* 32-bit package has been removed in Red Hat Enterprise Linux 7.2 and was replaced with a multilib version. If you have the 32-bit version of the package installed on a Red Hat Enterprise Linux 7.1 system, then you will encounter a **yum** error when attempting to upgrade to Red Hat Enterprise Linux 7.2.

To work around this problem, use the **yum remove pygobject3-devel.i686** command as **root** to uninstall the 32-bit version of the package before upgrading your system.

Build requirements not defined correctly for Emacs

The *binutils* package earlier than version 2.23.52.0.1-54 causes a segmentation fault during the build. As a consequence, it is not possible to build the Emacs text editor on IBM Power Systems. To work around this problem, install the latest *binutils*.

External display issues when combining laptop un/dock and suspend

In the GNOME desktop environment, with some laptops, external displays connected to a docking station might not be automatically activated when resuming a suspended laptop after it has been undocked and docked again.

To work around this problem, open the Displays configuration panel or run the **xrandr** command in a terminal. This makes the external displays available again.

Emacs sometimes terminates unexpectedly when using the up arrow on ARM

On the ARM architecture, the **Emacs** text editor sometimes terminates unexpectedly with a segmentation fault when scrolling up a file buffer. This happens only when the syntax highlighting is enabled. There is not currently any known workaround for this problem.

Chapter 35. Installation and Booting

Installation fails with a traceback when specifying %packages --nobase --nocore in a Kickstart file

Using a Kickstart file which contains the **%packages** section and specifies the **--nobase** and **--nocore** options at the same time causes the installation to fail with a traceback message due to the *yum-langpacks* package missing.

To work around this problem, add the *yum-langpacks* package within the **%packages** section when using **%packages --nobase --nocore** in your Kickstart file.

Installation can not proceed if a root password specified in Kickstart does not pass policy requirements

If you use a Kickstart file that defines a root password and the password does not fullfill requirements for the security policy selected in the Security Policy spoke, you will be unable to complete the installation. The **Begin Installation** button will be grayed out, and it is not possible to change the root password manually before pressing this button.

To work around this problem, make sure that your Kickstart file uses a sufficiently strong password that passes requirements defined by the selected security policy.

Rescue mode fails to detect and mount root volume on Btrfs

The installer rescue mode (accessed from the installation media boot menu or using the **inst.rescue** boot option) can not detect an existing system with the **/** (root) directory placed on a Btrfs subvolume. Instead, an error message saying 'You don't have any linux partitions.' is displayed.

To work around this problem, enter the shell and mount the root volume manually.

Wrong window title in Initial Setup

The Initial Setup tool, which is automatically displayed after the first post-installation reboot and which allows you to configure settings like network connections and to register your system, displays the string **__main__.py** in the window title.

This is a cosmetic problem and has no negative impact on usability.

Reinstalling on an FBA DASD on IBM System z causes the installer to crash

When reinstalling Red Hat Enterprise Linux 7 on IBM System z with a Fixed Block Architecture (FBA) DASD, the installer will crash due to incomplete support for these devices.

To work around this problem, ensure that any FBA DASDs are not present during the installation by placing them on the device ignore list. This should be done before launching the installer. From a root shell, use the **chccwdev** command followed by the **cio_ignore** command to manually switch devices offline and then add them to the device ignore list.

Alternatively, you can remove all FBA DASD device IDs from the CMS configuration file or the parameter file instead of using these commands before beginning the installation.

HyperPAV aliases are not available after installation on IBM System z

A known issue prevents DASDs configured as HyperPAV aliases from being automatically attached to the system after the installation finishes. These storage devices are available at the Installation Destination screen during installation, but they are not immediately accessible after you finish installing and reboot.

To fix this problem temporarily (until the next reboot), remove these devices from the device blacklist using the **chccwdev** command:

```
# chccwdev -e <devnumber>
```

To make the HyperPAV aliases available persistently across reboots, add their device numbers into the **/etc/dasd.conf** configuration file.

You can use the **lsdasd** command to verify that these devices are available.

Generated anaconda-ks.cfg file on IBM System z can not be used to reinstall the system

The **anaconda-ks.cfg** file, which is a Kickstart file generated during system installation and which contains all selections made during the install process, represents disk sizes as decimal numbers on IBM System z DASDs. This is because DASDs report a 4KiB alignment, which makes the calculated disk sizes incorrect as they are recorded in the Kickstart file, since only integer values are accepted. Therefore, it is not possible to re-use the generated Kickstart file to reproduce the installation.

Using the **anaconda-ks.cfg** file on IBM System z to reinstall the system requires you to manually change all decimal values within to integers.

Possible NetworkManager error message during installation

When installing the system, the following error message can be displayed and logged:

ERR NetworkManager: <error> [devices/nm-device.c:2590] activation_source_schedule(): (eth0): activation stage already scheduled

The error message should not prevent the installation from completing.

Package libocrdma is missing from the InfiniBand Support package group

The *libocrdma* package is not included in the default package set of the InfiniBand Support group. Consequently, when users select the InfiniBand Support group and are expecting RDMA over Converged Ethernet (RoCE) to work on Emulex OneConnect adapters, the necessary driver, *libocrdma*, is not installed by default.

On first boot, the user can manually install the missing package by issuing this command:

```
# yum install libocrdma
```

Alternatively, add the *libocrdma* package to the **%packages** section of your Kickstart file.

As a result, the user will now be able to use the Emulex OneConnect devices in RoCE mode.

Insufficient size of the /boot partition may prevent the system from upgrading

The /boot partition, which contains installed kernels and initial ram disks, may become full if multiple kernels and additional packages such as *kernel-debug* are installed. This is caused by the default size of this partition being set to 500 MB during installation, and prevents the system from being upgraded.

As a workaround, use **yum** to remove older kernels if you do not need them. If you are installing a new system, you should also consider this possibility, and set the /boot partition to a larger size (for example 1 GB) instead of the default (500 MB).

Installation on multipath devices fails if one or more disks are missing a label

When installing on multipath devices, the installer may display an error dialog if it fails to read one or more disks which are a member of the multipath. This problem is caused by one or more disks missing a disk label, and the installation can not proceed if it occurs.

To work around this problem, create disk labels on all disks which are part of the multipath device you are using during the installation.

Static IPv4 configuration in Kickstart is overwritten if a host name is defined in %pre script

When defining a host name in the **%pre** section of a Kickstart file, a **network** command that only sets host name ("network --hostname=hn") is considered as a device configuration with default --bootproto value ("dhcp") and default **--device** value ("link", which means the first device with link found). The Kickstart then behaves as if **network --hostname=hn --device=link** was used.

If the device considered as default for the **--device** option (the first device with link found) has already been configured to use static IPv4 configuration (for example with the preceding **network** command), the configuration is overriden by the IPv4 DHCP implied by the **--hostname** option.

To work around this problem, make sure that the **network** command which defines the host name is used first, and the second **network** command which would normally be overridden is used afterwards.

In cases where the **network** command defining a host name is the only such command in the Kickstart file, add a **--device** option to it with a non-existing interface (for example, **network --hostname=hn --device=x**).

Using the realm command in Kickstart causes the installer to crash

A known issue prevents the **realm** command from being used in Kickstart files. Attempting to join an Active Directory or Identity Management domain during the installation using this command causes the installer to crash.

To work around this problem, you can either wait until the installation finishes and join a domain manually afterwards, or you can add the **realm join <realm name>** command to the Kickstart file's **%post** section. See the **realm(8)** man page for information joining a domain using the command line.

Installer built-in help is not updated during system upgrade

When upgrading from Red Hat Enterprise Linux 7.1 to version 7.2, the built-in help for the Anaconda installer (the *anaconda-user-help* package) is not upgraded due to a significant change in packaging.

To work around this problem, use **yum** to remove the *anaconda-user-help* package before performing the upgrade, and install it again after you finish upgrading to Red Hat Enterprise Linux 7.2.

Incorrect ordering of boot menu entries generated by grubby

The **grubby** tool, which is used to modify and update the GRUB2 boot loader configuration files, may add debug boot menu entries at the top of the list when generating the boot menu configuration file. These debug menu entries then cause normal entries to be pushed down, although they are still highlighted and selected by default.

Using multiple driver update images at the same time only applies the last one specified

When attempting to perform a driver update during the installation using the **inst.dd=/dd.img** boot option and specifying it more than once to load multiple driver update images, Anaconda will ignore all instances of the parameter except the last one.

To work around this problem, you can:

* Install additional drivers after the installation if possible

* Use alternate means to specify a driver update image, such as the **driverdisk** Kickstart command

* Combine multiple driver update images into a single one

Installer crashes when it detects LDL-formatted DASDs

The installer crashes whenever it detects the LDL (Linux Disk Layout) format on one or more DASDs on IBM System z. The crash is caused by a race condition in the **libparted** library and happens even if these DASDs are not selected as installation targets. Other architectures are not affected by this issue.

If LDL DASDs are to be used during installation, users should manually reformat each LDL DASD as CDL (Compatible Disk Layout) using the **dasdfmt** command in a root shell before launching the installer.

If LDL DASDs are present on a system and a user does not wish to utilize them during installation, they should be placed on the device ignore list for the duration of the installation process. This should be done before launching the installer. From a root shell, users should use the **chccwdev** command followed by the **cio_ignore** command to manually switch devices offline and then add them to the device ignore list.

Alternatively, you can remove all LDL DASD device IDs from the CMS configuration file or the parameter file instead of using these commands before beginning the installation.

Chapter 36. Kernel

Some ext4 file systems cannot be resized

Due to a bug in the ext4 code, it is currently impossible to resize ext4 file systems that have 1 kilobyte block size and are smaller than 32 megabytes.

Repeated connection loss with iSER-enabled iSCSI targets

When using the server as an iSER-enabled iSCSI target, connection losses occur repeatedly, the target can stop responding and the kernel becomes unresponsive. To work around this issue, minimize iSER connection losses or revert to non-iSER iSCSI mode.

Installer does not detect Fibre Channel over Ethernet disks on EDD systems

On EDD systems, FCoE disks are not detected automatically by Anaconda due to the **edd** driver missing. This makes such disks unusable during the installation.

To work around this problem, perform the following steps:

* Add **fcoe=edd:nodcb** to the kernel command line during the installation, the FCoE disks will be detected by anaconda.

* Add **fcoe=edd:nodcb** to the rescue image and boot the system with it.

* Add the edd module to the initrd image by executing the following commands:

```
#dracut --regenerate-all -f

#dracut --add-drivers edd /boot/initramfs-3.10.0-123.el7.x86_64.img
```

* Reboot the system with the default boot menu entry

NUMA balancing does not work optimally under certain circumstances

The Linux kernel Non-Uniform Memory Access (NUMA) balancing does not work optimally under the following condition in Red Hat Enterprise Linux 7. When the **numa_balancing** option is set, some of the memory can move to an arbitrary non-destination node before moving to the constrained nodes, and the memory on the destination node also decreases under certain circumstances. There is currently no known workaround available.

PSM2 MTL disabled to avoid conflicts between PSM and PSM2 APIs

The new *libpsm2* package provides the PSM2 API for use with Intel Omni-Path devices, which overlaps with the Performance Scaled Messaging (PSM) API installed by the *infinipath-psm* package for use with Truescale devices. The API overlap results in undefined behavior when a process links to libraries provided by both packages. This problem affects **Open MPI** if the set of its enabled MCA modules includes the **psm2** Matching Transport Layer (MTL) and one or more modules that directly or indirectly depend on the **libpsm_infinipath.so.1** library from the *infinipath-psm* package.

To avoid the PSM and PSM2 API conflict, Open MPI's **psm2** MTL has been disabled by default in the **/etc/openmpi-*/openmpi-mca-params.conf** configuration file. If you enable it, you need to disable the **psm** and **ofi** MTLs and the **usnic** Byte Transfer Layer (BTL) that conflict with it (instructions are also provided in comments in the configuration file).

There is also a packaging conflict between the *libpsm2-compat-devel* and *infinipath-psm-devel*

packages because they both contain PSM header files. Therefore, the two packages cannot be installed at the same time. To install one, uninstall the other.

Performance problem of the perf utility

The **perf archive** command, which creates archives with object files with build IDs found in **perf.data** files, takes a long time to complete on IBM System z. At present, no known workaround exists. Other architectures are not affected.

qlcnic fails to enslaved by bonding

Certain bonding modes set a MAC address on the device which the qlcnic driver does not properly recognize. This prevents the device from restoring its original MAC address when it is removed from the bond.

As a workaround, unenslave the qlcnic driver and reboot your operating system.

Installation fails on some 64-bit ARM Applied Micro computers

Red Hat Enterprise Linux 7.2 fails to install on certain 64-bit ARM systems by Applied Micro with the following error message:

Unable to handle kernel NULL pointer dereference at virtual address 0000033f

At present, there is no workaround for this problem.

libvirt management of VFIO devices can lead to host crashes

The **libvirt** management of host PCI devices, assigned to guests using the VFIO driver, can lead to host kernel drivers and the vfio-pci driver binding simultaneously to devices in the same IOMMU group. This is an invalid state, which can lead to a host unexpected termination.

For now, the only workaround is to never hot-unplug a VFIO device from a guest, if there are any other devices in the same IOMMU group.

Installation using iSCSI and IPv6 hangs for 15 minutes

Dracut times out after trying to connect to the specified iSCSI server for 15 minutes if IPv6 is enabled. Eventually, Dracut connects successfully and proceeds as expected; however, to avoid the delay, use **ip=eth0:auto6** on the installer's command line.

i40e NIC freeze

With old firmware, a network card using the i40e driver becomes unusable for about ten seconds after it enters the promiscuous mode. To avoid this problem, update the firmware.

i40e is issuing WARN_ON

The i40e driver is issuing the WARN_ON macro during ring size changes because the code is cloning the rx_ring struct but not zeroing out the pointers before allocating new memory. Note that this warning is harmless to your system.

neprio_cgroups not mounted at boot

Currently, systemd mounts the /sys/fs/cgroup/ directory as read-only, which prevents the default mount of the /sys/fs/cgroup/net_prio/ directory. As a consequence, the netprio_cgroups module is not mounted at boot. To work around this problem, use the **`mount -o remount`** command, followed by **`rw -t cgroup nodev /sys/fs/cgroups`**. This makes it possible to install module-based cgroups manually.

Chapter 37. Networking

Timeout policy not enabled in Red Hat Enterprise Linux 7.2 kernel

The **`nfct timeout`** command is not supported in Red Hat Enterprise Linux 7.2. As a workaround, use the global timeout values available at **`/proc/sys/net/netfilter/nf_conntrack_*_timeout_*`** to set the timeout value.

Chapter 38. Storage

No support for thin provisioning on top of RAID in a cluster

While RAID logical volumes and thinly provisioned logical volumes can be used in a cluster when activated exclusively, there is currently no support for thin provisioning on top of RAID in a cluster. This is the case even if the combination is activated exclusively. Currently this combination is only supported in LVM's single machine non-clustered mode.

When using thin-provisioning, it is possible to lose buffered writes to the thin-pool if it reaches capacity

If a thin-pool is filled to capacity, it may be possible to lose some writes even if the pool is being grown at that time. This is because a resize operation (even an automated one) will attempt to flush outstanding I/O to the storage device prior to the resize being performed. Since there is no room in the thin-pool, the I/O operations must be errored first to allow the grow to succeed. Once the thin pool has grown, the logical volumes associated with the thin-pool will return to normal operation.

As a workaround to this problem, set 'thin_pool_autoextend_threshold' and 'thin_pool_autoextend_percent' appropriately for your needs in the lvm.conf file. Do not set the threshold so high or the percent so low that your thin-pool will reach full capacity so quickly that it does not allow enough time for it to be auto-extended (or manually extended if you prefer). If you are not using over-provisioning (creating logical volumes in excess of the size of the backing thin-pool), then be prepared to remove snapshots as necessary if the thin-pool begins to near capacity.

Chapter 39. System and Subscription Management

Non-working `Back` button in the Subscription Manager add-on for Initial Setup

The **Back** button on the first panel of the Subscription Manager add-on for the Initial Setup utility does not work. To work around this problem, click **Done** at the top of Initial Setup to exit the registration workflow.

virt-who fails to change host-to-guest association to the Candlepin server

When adding, removing, or migrating a guest, the virt-who utility currently fails to send the host-to-guest mapping and prints a RateLimitExceededException error to the log file. To work around the problem, set the **VIRTWHO_INTERVAL=** parameter in the /etc/sysconfig/virt-who file to a large number, such as 600. This allows the mapping to be changed correctly, but causes changes in the host-to-guest mapping to take significantly longer to be processed.

Chapter 40. Virtualization

Problematic GRUB 2 navigation with KVM

When using the serial console through KVM, holding down an arrow key for an extended period of time to navigate in the GRUB 2 menu results in erratic behavior. To work around this problem, avoid the rapid input caused by holding an arrow key down for a longer time.

Resizing GUID Partition Table (GPT) disks on Hyper-V guests causes partition table errors

The Hyper-V manager supports shrinking a GPT-partitioned disk on a guest if there is free space after the last partition, by allowing the user to drop the unused last part of the disk. However, this operation will silently delete the secondary GPT header on the disk, which may trigger error messages when guest examines the partition table (for example, with parted(8)). This is a known limit of Hyper-V.

To work around this, it is possible to manually restore the secondary GPT header with the gdisk(8) expert command **e**, after shrinking the GPT disk. This also occurs when using Hyper-V's Expand option, but can also be fixed with the parted(8) tool.

Bridge creation with `virsh iface-bridge` fails

When installing Red Hat Enterprise Linux 7 from other sources than the network, network device names are not specified by default in the interface configuration files (this is done with a **`DEVICE=`** line). As a consequence, creating a network bridge by using the **`virsh iface-bridge`** command fails with an error message. To work around the problem, add **`DEVICE=`** lines into the /etc/sysconfig/network-scripts/ifcfg-* files.

QEMU-emulated CAC smart cards incompatible with ActivClient software

Currently, Common Access Card (CAC) smart cards emulated with QEMU are not accepted by ActivClient software. To work around this problem, disable the pcscd daemon, provision a Windows KVM guest, preconfigure it in the virt-viewer tool and select the USB redirection option, install the ActivClient software, and reboot the KVM guest. With this setup, ActivClient accepts the emulated CAC card.

Chapter 41. Atomic Host and Containers

Atomic Host installation offers cryptsetup although it is not available

During the installation of Red Hat Enterprise Linux 7 Atomic Host, the installer offers the option to encrypt partitions using **cryptsetup** in the Manual Partitioning screen, the same way it offers to do so during Red Hat Enterprise Linux 7.2 installation.

However, encrypted partitions are not supported on Atomic Host. If you encrypt any partitions during the installation, you will not be able to unlock them later.

To work around this problem, do not encrypt any partitions or logical volumes when installing Red Hat Enterprise Linux Atomic Host, even though the installer presents this option.

Installer can only add advanced storage the first time the storage spoke is entered

During an interactive installation using the Anaconda graphical interface, adding advanced storage (iSCSI, zFCP, FCoE) to your disk selection does not work if you have already entered and left the storage spoke. To work around the problem, make sure the network, if needed, is active then enter the storage spoke and add all advanced storage devices.

Atomic Host installation offers BTRFS but it is not supported

The Red Hat Enterprise Linux Atomic Host installer offers BTRFS as a partition option, but the tree does not include btrfs-progs. Consequently, an Atomic Host system with BTRFS partitioning will not work even though the option present in the installer. Do not choose this option in the installer. BTRFS is not currently supported for Atomic Host.

ostreesetup in Kickstart files supports only HTTP and HTTPs

The Atomic Host **ostreesetup** Kickstart command only supports HTTP and HTTPS Uniform Resource Identifiers (URIs). Providing a different option, for example, **ftp://** can cause the installer to crash. Use only HTTP or HTTPS.

Customization of the host system not supported

Red Hat Enterprise Linux Atomic Host does not include a mechanism to customize or override the content of the host itself, for example it does not include a tool to use a custom kernel for debugging.

Red Hat Enterprise Linux Atomic Host only supports the en_US.UTF-8 locale

During installation, if you select a language other than American English as the default keyboard type, it is not reflected in the installed system afterwards. The locale is still set to **en_US** and error messages about locales missing are displayed. This could be problematic for programs that require other locales, or, for example, when you have a password in another language, the system will not recognize it.

When the root partition runs out of free space

Red Hat Enterprise Linux Atomic Host allocates 3GB of storage to the root partition, which includes the docker volumes (units of storage that a running container can request from the host system). This makes it easy for the root partition to run out of storage space. In order to support more volume space, more physical storage must be added to the system, or the root Logical Volume must be extended.

By default, 40% from the other volume, will be reserved for storing the container images. The other 60% can be used to extend the root partition. For detailed instructions, see https://access.redhat.com/documentation/en/red-hat-enterprise-linux-atomic-host/version-7/getting-started-with-containers/#changing_the_size_of_the_root_partition_after_installation.

Rescue mode does not work in Red Hat Enterprise Linux Atomic Host

The Anaconda installer is unable to find a previously installed Atomic Host system when in rescue mode. Consequently, rescue mode does not work and should not be used.

The docker daemon is unable to create a core dump

On Red Hat Enterprise Linux Atomic Host, the core pattern is set to **`core`**. This prevents writing core dumps for daemons like docker whose directory is **`root`** (**`/`**), because it is read-only. To work around this problem, specify a core file name pattern that points to a writeable location:

```
echo /var/lib/core > /proc/sys/kernel/core_pattern
```

With this workaround, core dumps will be saved under **`/var/lib`**.

The `brandbot.path` service may cause subscription-manager to change the /etc/os-release file in 7.1 installations

The **`/etc/os-release`** file may still specify the 7.1 version even after Atomic Host has been upgraded to 7.2 using the **`atomic host upgrade`** command. This occurs because the underlying ostree tool preserves modified files in **`/etc`**. As a workaround, after upgrading to 7.2, run the following command: **`cp /usr/etc/os-release /etc`**. This way, the **`/etc/os-release`** file will return to an unmodified state, and because *brandbot.path* is masked in 7.2.0, it will not be modified in the future by subscription-manager, and future upgrades will show the correct version.

Appendix A. Component Versions

This appendix is a list of components and their versions in the Red Hat Enterprise Linux 7.2 release.

Table A.1. Component Versions

Component	Version
Kernel	3.10.0-327
QLogic **qla2xxx** driver	8.07.00.08.07.2-k
QLogic **qla4xxx** driver	5.04.00.00.07.02-k0
Emulex **lpfc** driver	0:10.7.0.1
iSCSI initiator utils	*iscsi-initiator-utils-6.2.0.873-32*
DM-Multipath	*device-mapper-multipath-0.4.9-85*
LVM	*lvm2-2.02.130-5*

Appendix B. Revision History

Revision 0.0-1.20	**Thu Nov 19 2015**	**Lenka Špačková**

Release of the Red Hat Enterprise Linux 7.2 Release Notes.

Revision 0.0-1.4	**Mon Aug 31 2015**	**Laura Bailey**

Release of the Red Hat Enterprise Linux 7.2 Beta Release Notes.