Red Hat Enterprise Linux 6

6.9 Technical Notes

Technical Notes for Red Hat Enterprise Linux 6.9

Edition 9

Last Updated: 2017-10-27

Red Hat Enterprise Linux 6 6.9 Technical Notes

Technical Notes for Red Hat Enterprise Linux 6.9
Edition 9

Red Hat Customer Content Services
rhel-notes@redhat.com

Legal Notice

Abstract

The Technical Notes provide information about notable bug fixes, Technology Previews, deprecated functionality, and other details in Red Hat Enterprise Linux 6.9. For high-level coverage of the improvements implemented in Red Hat Enterprise Linux 6.9 and a list of known problems in this release, refer to the Release Notes. TODO: update link for Beta/GA

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 6.9 Technical Notes*document provides a list of notable bug fixes, all currently available Technology Previews, deprecated functionality, and other information. The Release Notes document describes the major changes made to the Red Hat Enterprise Linux 6 operating system and its accompanying applications for this minor release, as well as known problems.

Capabilities and limits of Red Hat Enterprise Linux 6 as compared to other versions of the system are available in the Red Hat Knowledgebase 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. RED HAT ENTERPRISE LINUX 6.9 INTERNATIONAL LANGUAGES

Red Hat Enterprise Linux 6.9 supports installation of multiple languages and changing of languages based on your requirements.

The following languages are supported in Red Hat Enterprise Linux 6.9:

- East Asian Languages - Japanese, Korean, Simplified Chinese, and Traditional Chinese
- European Languages - English, German, Spanish, French, Portuguese Brazilian, and Russian,

The table below summarizes the currently supported languages, their locales, default fonts installed and packages required for some of the supported languages

Table 1.1. Red Hat Enterprise Linux 6 International Languages

Territory	Language	Locale	Fonts	Package Names
China	Simplified Chinese	zh_CN.UTF-8	AR PL (ShanHeiSun and Zenkai) Uni	fonts-chinese, scim-pinyin, scim-tables
Japan	Japanese	ja_JP.UTF-8	Sazanami (Gothic and Mincho)	fonts-japanese, scim-anthy
Korea	Hangul	ko_KR.UTF-8	Baekmuk (Batang, Dotum, Gulim, Headline)	fonts-korean, scim-hangul
Taiwan	Traditional Chinese	zh_TW.UTF-8	AR PL (ShanHeiSun and Zenkai) Uni	fonts-chinese, scim-chewing, scim-tables
Brazil	Portuguese	pt_BR.UTF-8	standard latin fonts	
France	French	ft_FR.UTF-8	standard latin fonts	
Germany	German	de_DE.UTF-8	standard latin fonts	
Italy	Italy	it_IT.UTF-8	standard latin fonts	
Russia	Russian	ru_RU.UTF-8	Cyrillic	dejavu-lgc-sans-fonts, dejavu-lgc-sans-mono-fonts, dejavu-lgc-serif-fonts, xorg-x11-fonts-cyrillic

Territory	Language	Locale	Fonts	Package Names
Spain	Spanish	es_ES.UTF-8	standard latin fonts	

PART I. NOTABLE BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 6.9 that have a significant impact on users.

CHAPTER 2. AUTHENTICATION AND INTEROPERABILITY

SSSD correctly reports supplementary groups for AD users in a nested domain

Resolving supplementary groups sometimes failed for Active Directory (AD) users with the same **samAccountName** attribute who existed in two AD domains, when:

- one of the AD domains was nested under the other
- the users were stored in a non-default organizational unit (OU)

Consequently, the **id [user_name]** command reported only the primary group for these users.

The underlying SSSD code has been improved to better match the user account with its domain. As a result, SSSD reports also supplementary groups of AD users in the described situation. (BZ#1293168)

Authentication no longer fails when two SRV resolution requests are running at the same time

When multiple service record (SRV) resolution requests were running concurrently, one of them returned a failure indicating that no new servers were found. Consequently, authentication using the **ssh** utility failed. With this update, SSSD handles two concurrent SRV resolution requests gracefully. As a result, authentication no longer fails in this situation. (BZ#1367435)

Users with expired or locked accounts now cannot log in to IdM clients with their SSH keys

When a trusted Active Directory (AD) user with an expired or locked user account attempted to log in to an Identity Management (IdM) client using a non-password login method, such as SSH keys, the login was successful. With this update, the IdM client checks the AD lockout attribute when verifying whether an AD user is allowed to log in. As a result, AD users with expired or locked accounts are no longer permitted to log in in this situation.

Note that this bug has no security impact: The AD user could not obtain a Kerberos ticket on the IdM client because the user account was expired or locked on the server side. (BZ#1335400)

sssd_be subprocesses no longer unnecessarily consume memory

Previously, when the **id_provider** option was set to **ad** in the **/etc/sssd/sssd.conf** file, a helper process inside the **sssd_be** process sometimes failed. In consequence, the process was spawning new **sssd_be** instances, which consumed additional memory.

With this update, SSSD does not fork **sssd_be** subprocesses if no helper program is available. This reduces the amount of consumed memory. (BZ#1336453)

Attempts to renew the system password in a keytab no longer cause SSSD to stop working

When attempting to renew the system password stored in a keytab, System Security Services Daemon (SSSD) leaked a file descriptor. The leaked file descriptors gradually accumulated, which caused SSSD to stop working.

With this update, SSSD no longer leaks file descriptors in this situation. As a result, SSSD is able to keep updating the system password without the described negative impact on the system. (BZ#1340176)

SSSD now correctly processes GPO files that contain attributes in a format other than key=value

Previously, System Security Services Daemon (SSSD) did not correctly process INI files that contained attribute pairs in a format other than **key=value**. Consequently, SSSD failed to process group policy

object (GPO) files that contained such attributes.

This update ensures that SSSD processes the mentioned files correctly even if they use a different attribute format than **key=value**. (BZ#1374813)

SSSD now resolves users with `externalUser` correctly

Support for the **externalUser** LDAP attribute was removed from the System Security Services Daemon (SSSD) in Red Hat Enterprise Linux 6.8. In consequence, the assignment of **sudo** rules to local accounts, such as by using the **/etc/passwd** file, failed. The problem affected only accounts outside of Identity Management (IdM) domains and Active Directory (AD) trusted domains.

This update ensures that SSSD correctly resolves users with the **externalUser** attribute defined. As a result, assigning **sudo** rules works as expected in the described situation. (BZ# 1321884)

SSSD correctly creates local overrides in an AD environment

Previously, the **sss_override** utility created case-insensitive distinguished names (DNs) when the **id_provider** option was set to **ad** in the **/etc/sssd/sssd.conf** file. However, the DNs in the SSSD cache are stored as case-sensitive. As a consequence, local overrides were not created for users from the Active Directory (AD) subdomain and for users with mixed-case account names. With this update, SSSD searches the object in the cache and uses the DN from the search result. This fixes the problem in the mentioned situation. (BZ#1327272)

OpenLDAP now correctly sets NSS settings

Previously, the OpenLDAP server used an incorrect handling of network security settings (NSS) code. As a consequence, settings were not applied, which caused certain NSS options, such as **olcTLSProtocolMin**, not to work correctly. This update addresses the bug and as a result, the affected NSS options now work as expected. (BZ#1249092)

IPA replica installation no longer fails due to malformed HTTP requests

A bug in pki-core previously caused PKI to generate HTTP requests missing a **Host** header and using incorrect line delimiters during IPA replica installation. At the same time, an update to **httpd** caused these malformed requests to be rejected, even though they were accepted in previous versions, and as a result, IPA replica installations failed. This update to pki-core fixes the problem in HTTP request generation, and replica installations now work as expected. (BZ#1403943)

CHAPTER 3. CLUSTERING

The PCS cluster stop operation now completes successfully when cluster nodes include resources that require DLM

When stopping the cluster on all nodes by running **pcs cluster stop --all**, resources that require the Distributed Lock Manager (DLM), such as gfs2 or clustered logical volumes, in some cases lost quorum before they were able to shut down. As a consequence, the stop operation became unresponsive. With this update, **pcs cluster stop --all** stops the cman service on all nodes only after Pacemaker has stopped those nodes. As a result, quorum is maintained while all resources are stopping, and the operation is thus able to complete successfully. (BZ#1322595, BZ#1353738)

The rgmanager daemon can now correctly start clustered services on surviving nodes when quorum is regained

With central processing mode enabled, when quorum was dissolved and regained, the rgmanager daemon stopped working on a surviving cluster node. With this update, the configuration tree is repopulated after quorum is regained. As a result, clustered services start up on the surviving cluster node as expected in the described scenario. (BZ#1084053)

Short time between the start of **rgmanager** and **clustat** no longer leads to **rgmanager** crashing

When the **clustat** utility was run shortly after the **rgmanager** daemon started but before it completely finished initializing, **rgmanager** was susceptible to unexpected termination. This bug has been fixed and **rgmanager** now starts without crashing in this scenario. (BZ#1228170)

rgmanager exits without problems after **cman** is stopped

When the **cman** service was stopped before the **rgmanager** daemon, **rgmanager** in some cases exited unexpectedly on cluster nodes. With this update, the **cpg_lock()** function has been fixed and **rgmanager** exits gracefully in the described scenario. (BZ#1342825)

Time-related values of cluster resource configuration are now evaluated properly

Previously, time-related resource values in actual use could differ from the values configured in the **cluster.conf** file, especially at the initial configuration load. This could cause the **rgmanager** daemon to behave unpredictably. With this fix, **rgmanager** behaves exactly as configured with regards to resources and respective time-related values. (BZ#1414139)

CHAPTER 4. COMPILER AND TOOLS

Resolution for gcc compatibility issue with `sockaddr_in`

The default RHEL 6 compiler has rules about how to copy **`struct sockaddr_in`** that are substantially different to those used by later versions of gcc (E.g. the tools provided with Red Hat Developer Toolset). This caused corrupted sockaddr_in overlays when using newer compilers.

This update changes the way the **unused** portions of the sockaddr_in structure are defined. Now, newer versions of gcc will copy them correctly.

Note that -fno-strict-aliasing is still required for compiling such sources. (BZ#1338673)

Resolution for floating point exception error when measuring memory usage of processes that did not allocate memory

When running the memusage utility on programs that did not explicitly allocate any further memory, a floating point exception was encountered.

This update checks for zeroed internal statistics and will not divide by them, avoiding the exception. (BZ#1331304)

Improved behavior in getaddrinfo() when scanning interfaces after being passed an IP address

Versions of glibc prior to this one would scan every IP address on an interface referenced by a getaddrinfo() call, even if the hostname passed was itself a numeric IP address. On systems with large numbers of IP addresses configured on the interface, this caused the call to take an excessive amount of time.

With this update, the scan happens only when needed and the call returns quickly when passed a numeric IP address. (BZ#1270950)

Fix for handling any open file descriptors in the event of thread cancellation

The use of POSIX thread cancellation could cause glibc to improperly handle open file descriptors, particularly those held open when processing identity information.

To correct this and ensure that functions like getpwuid_r complete, even when the thread is being cancelled, the library calls have been changed to correctly handle open file descriptors in any call from the exec family of functions. (BZ#1012343)

Fix for tzdata-update inheriting an unusable umask

When updating /etc/localtime, tzdata-update applies the current process umask to determine the file permissions.

If the umask is a restrictive value, such as 077, the new /etc/localtime file may not be readable by non-root users. To resolve this, tzdata-update now sets the permissions to rw-r--r-- (0644) unconditionally. (BZ#1373646)

Resolution for getaddrinfo accessing uninitialised data

On systems with nscd enabled, the getaddrinfo function in glibc could access uninitialized data and return false address information.

This update avoids accessing uninitialized data and ensures that correct addresses are returned. (BZ#1223095)

The system default CA bundle has been set as default in the compiled-in default setting or configuration in `mutt`

Previously, when connecting to a new system via TLS/SSL, the **mutt** email client required the user to save the certificate. With this update, the system Certificate Authority (CA) bundle is set in **mutt** by default. As a result, **mutt** now connects via SSL/TLS to hosts with a valid certificate without prompting the user to approve or reject the certificate. (BZ#1196787)

Resolv::DNS no longer returns truncated DNS replies

The Ruby **Resolv::DNS** resolver silently returned truncated DNS replies when the DNS response did not fit into the hard-coded 512-byte limit. Consequently, numerous DNS records required upgrading of the DNS connection from User Datagram Protocol (UDP) to Transmission Control Protocol (TCP) to receive complete DNS replies. This affected, for example, Microsoft Active Directory domains where the Key Distribution Center (KDC) list contained a larger set of servers. With this update, if a UDP reply is truncated, **Resolv::DNS** retries using TCP, which is the correct behavior according to RFC 1123. As a result, complete DNS replies are returned. (BZ#1331086)

tcsh no longer becomes unresponsive when the .history file is located on a network file system

Previously, if the **.history** file was located on a network file system, such as NFS or Samba, the **tcsh** command language interpreter sometimes became unresponsive during the login process. With this update, the **.history** file is not locked if located on a network file system. As a result, **tcsh** no longer becomes unresponsive in the described situation.

Note that having multiple instances of **tcsh** running can cause the **.history** file to become corrupted. You can resolve this problem by enabling explicit file-locking mechanism. To do that, add the **lock** parameter to the **savehist** option in the **tcsh** configuration file. For example:

```
$ cat /etc/csh.cshrc
# csh configuration for all shell invocations.
set savehist = (1024 merge lock)
```

To force **tcsh** to use file-locking when **.history** is located on a network file system, the **lock** parameter must be the third parameter of the **savehist** option. Do this at your own risk, because Red Hat does not guarantee that using the **lock** parameter prevents **tcsh** from becoming unresponsive during the login process. (BZ#885901)

The LWP::UserAgent Perl module now correctly handles proxy settings for HTTPS requests

The **LWP::UserAgent** Perl module previously did not honor HTTPS proxy environment variables by default. The perl-libwww-perl package version 5.883-3 started using the **IO::Socket::SSL** module instead of the **Net::SSL** module for implementing TLS. Consequently, applications that rely on processing of the **https_proxy** environment variable in the **Net:SSL** module established connections directly to the HTTPS server instead of through the HTTPS proxy server.

With this update, the **Net::SSL** module's behavior has been added to the **LWP::UserAgent** module to ensure that the **https_proxy** and **HTTPS_PROXY** environment variables are honored if no **env_proxy** option has been passed to the **LWP::UserAgent** module's **new()** method. Additionally, proxy specifications without a URL schema are now recognized. As a result, applications using the **Net::SSL** module correctly work after switching to the **IO::Socket::SSL** cryptographic back end from the perl-libwww-perl package. (BZ#1400632)

The Frontier::Client Perl module no longer ignores proxy settings for HTTPS requests

When using the **Frontier::Client** Perl module to send an XML-RPC request to an HTTPS server through a proxy service, the proxy setting was previously ignored. Consequently, the request was sent

directly to the HTTPS server and not through the proxy server. With this update, the **Frontier::Client** Perl module has been corrected to pass the proxy setting to an underlying **LWP::UserAgent** object for both HTTP and HTTPS schemata, and **Frontier::Client** now respects proxy setting also for HTTPS requests. (BZ#832390)

RPM verification no longer reports failures in the /var/account/pacct file

Previously, the RPM verification and compliance check reported failures, because the mode of the **/var/account/pacct** file was not set after installation. To fix this bug, the mode of **/var/account/pacct** is changed immediately after installation to 0600. As a result, the RPM verification check no longer reports failures in the **/var/account/pacct**, and the file is now accessible only by root. (BZ#1182317)

Output of jobs in tcsh is now correctly displayed to stdout

Previously, the output of the **jobs** built-in command was displayed to standard error instead of standard output. This bug has been fixed, and the output of **jobs** is now correctly displayed to stdout. (BZ#1338986)

Several regressions in the tcsh have been fixed

This update fixes several regressions in the **tcsh** command language interpreter:

- When browsing command history using the **Ctrl+P** or **Up Arrow** keys, backslashes are now shown correctly.
- A backslash now correctly escapes user-defined aliases; this fixes git-completion.
- The output of the built-in **time** command is now correctly captured when used with the built-in **setenv** command. (BZ#1334751)

git shortlog no longer crashes due to using freed memory

Previously, when email address entries differed only in case, the **.mailmap** feature of the **git shortlog** command did not replace a duplicate email entry with a strdup pointer, and freed memory was referenced. Consequently, **Git** terminated unexpectedly due to using already freed memory. A patch has been applied, which ensures that memory is freed before these entries are replaced, and **git shortlog** correctly uses only allocated memory. (BZ# 874659)

Perl interpreter no longer crashes when attempting to report Can't coerce HASH to string

When running an **XML::LibXSLT** Perl script in a **mod_perl** environment, the Perl interpreter could terminate unexpectedly with a segmentation fault while it tried to report the **Can't coerce HASH to string** runtime error. The code printing the error message has been modified to handle missing script line details correctly, and the Perl interpreter no longer crashes in this scenario. (BZ#1364206)

gdbserver now supports seamless debugging of processes from containers

Prior to this update, when **GDB** was executing inside a Super-Privileged Container (SPC) and attached to a process that was running in another container on Red Hat Enterprise Linux Atomic Host, **GDB** did not locate the binary images of the main executable or any shared libraries loaded by the process to be debugged.

As a consequence, **GDB** could have displayed error messages relating to files not being present, or being present but mismatched. Also, **GDB** may have seemed to attach correctly, but subsequent commands may have failed or displayed corrupted information.

In Red Hat Enterprise Linux 6.9, **gdbserver** has been extended for seamless support of debugging processes from containers. The Red Hat Enterprise Linux 6.9 version of **gdbserver** newly supports

the **qXfer:exec-file:read** and **vFile:setfs** packets. However, the Red Hat Enterprise Linux 6.9 version of **gdb** cannot use these packets. The Red Hat Developer Toolset 4.1 (or later) version of **gdb** is recommended for use with containers and with Red Hat Enterprise Linux 6.9 **gdbserver**. The Red Hat Developer Toolset version of **gdbserver** can be used as well.

Red Hat Enterprise Linux 6.9 **gdb** can now suggest using **gdbserver** when run with the **-p** parameter (or the **attach** command) and when, at the same time, it detects that the process being attached is from a container. Red Hat Enterprise Linux 6.9 **gdb** now also suggests the explicit use of the **file** command to specify the location of the process executable in the container being debugged. The **file** command does not need to be entered when the Red Hat Developer Toolset version of **gdb** is being used instead.

With this update, Red Hat Enterprise Linux 6.9 **gdbserver** provides seamless debugging of processes from containers together with Red Hat Developer Toolset 4.1 (or later) **gdb**. Additionally, Red Hat Enterprise Linux 6.9 **gdb** guides the user through the debugging of processes from containers when Red Hat Developer Toolset **gdb** is not available. (BZ#1316539)

CHAPTER 5. DESKTOP

Cancelling shutdown from a GUI session now switches to running session

Previously, the `policykit` credentials for shutting down the system were not obtained up front before logging out. Consequently, if the users cancelled the policykit authentication dialog, they were sent back to the login screen instead of the graphical user interface (GUI) session. With this update, the dialog for obtaining credentials appears up front before the logout starts. As a result, when the users cancel the authentication dialog, they are sent back to their session instead of the login screen. (BZ#1320245)

CHAPTER 6. DIRECTORY SERVER IN RED HAT ENTERPRISE LINUX

Directory Server no longer logs false positive error messages

Previously, in a Directory Server multi-master replication environment, the **`Failed to update RUV for unknown`** error message was logged multiple times when only the replica update vector (RUV) was updated without any change. This update fixes the problem and now Directory Server no longer logs the error message. (BZ#1266920)

In FIPS mode, the `slapd_pk11_getInternalKeySlot()` function is now used to retrieve the key slot for a token

The Red Hat Directory Server previously tried to retrieve the key slot from a fixed token name, when FIPS mode was enabled on the security database. However, the token name can change. If the key slot is not found, Directory Server is unable to decode the replication manager's password and replication sessions fail. To fix the problem, the **`slapd_pk11_getInternalKeySlot()`** function now uses FIPS mode to retrieve the current key slot. As a result, replication sessions using **SSL** or **STTARTTLS** no longer fail in the described situation. (BZ#1352109)

Directory Server now supports configuring weak DH parameters

The network security services (NSS) libraries, linked with the Red Hat Directory Server, require a minimum of 2048-bit Diffie-Hellman (DH) parameters. However, Java 1.6 and 1.7 supports only 1024-bit DH parameters. As a consequence, clients using these Java versions were unable to connect to Directory Server using encrypted connections. This update adds the **`allowWeakDHParam`** parameter to the **`cn=encryption,cn=config`** entry. As a result, if this parameter is enabled, affected clients can now connect using weak DH parameters. (BZ#1327065)

The `cleanAllRUV` task no longer corrupts changelog back ends

At the end of the **`cleanAllRUV`** task, Directory Server removes entries from the replication changelog that contain the cleaned replica ID. Previously, the task incorrectly ran all changelog back ends instead of only the one set in the task. As a consequence, if multiple back ends contained the same replica ID, the **`cleanAllRUV`** task corrupted them. This update fixes the problem and now the **`cleanAllRUV`** task works correctly. (BZ#1369572)

Reindexing the `retro` changelog no longer fails

Previously, the **`retrocl-plugin`** sets a lock in read mode on the changelog back end without releasing it. This could result in a deadlock situation. For example, an index task executed by the **`db2index.pl`** script on the **`retro`** changelog back end became unresponsive when a lock in write mode was set. This update applies a patch and as a result, reindexing the **`retro`** changelog no longer fails. (BZ#1370145)

Directory Server no longer fails when disabling the `CLEAR` password storage scheme plug-in

Previously, Directory Server required that the **CLEAR** password storage plug-in was enabled when setting **`userPassword`** attributes. As a consequence, Directory Server terminated unexpectedly when attempting to set **`userPassword`** attributes, if **CLEAR** was disabled. This update applies a patch and as a result, Directory Server no longer fails in the described situation. (BZ#1371678)

Directory Server no longer terminates unexpectedly when using server side sorting

Previously, when using a matching rule and server side sorting, Directory Server incorrectly frees memory multiple times and terminates unexpectedly. This update fixes the bug, and as a result Directory Server no longer fails when using server side sorting. (BZ#1371706)

Directory Server now validates macros in ACIs

Previously, the Red Hat Directory Server did not validate macros in an access control instruction (ACI). As a result, users were able to set incorrect macros in an ACI. This update improves the code underlying validation, and Directory Server rejects invalid macros and logs an error. (BZ#1382386)

Replication monitor now shows the correct date

On the replication monitor, the year of the date was not displayed in the header when the value of the **day** field was less than 10. The code now uses the correct API, and the year is displayed correctly. (BZ#1410645)

The `memberOf` fix-up task now verifies arguments

Previously, if an invalid filter or **basedn** parameter was provided in the **`memberOf`** fix-up task, and the task failed, no information was logged. A patch has been applied and now, if a problem occurs, an error is logged and the task status is updated. As a result, the administrator is now able to identify if a task failed. (BZ#1406835)

Directory Server no longer terminates unexpectedly when deleting a non-existent attribute

Previously, deleting a non-existent attribute from the back end configuration caused Directory Server to terminate unexpectedly. This update applies a patch to pass a **`NULL`** value to the **`ldbm_config_set()`** function if no attribute was deleted. As a result, Directory Server now rejects the operation in the described scenario. (BZ#1403754)

Directory Server no longer displays multiple error messages when importing fails

Previously, if importing data failed, multiple **`Unable to flush`** error message were be displayed, because the connection to the database was not closed. This update applies a patch and as a result, Directory Server no longer displays multiple errors in the mentioned situation. (BZ#1402012)

Virtual list view-related problems have been fixed

Previously, when removing a virtual list view (VLV) index, the **`dblayer_erase_index_file_nolock()`** function was not called. Thus, the physical index file and the back pointer set to the **`dblayer`** handle were not removed. Consequently, Directory Server terminated unexpectedly. This fix updates the code and the **`dblayer_erase_index_file_nolock()`** function is now called when removing a VLV index.

In addition, the **`vlv_init()`** function previously could be called multiple times without unregistering VLV plug-in callbacks. As a consequence, Directory Server sometimes terminated unexpectedly. With this update, callbacks are now unregistered.

As a result, Directory Server no longer terminates unexpectedly in the described situations. (BZ#1399600)

Directory Server no longer logs sensitive information

Previously, when the **`Trace function calls`** option was enabled in the **`nsslapd-errorlog-level`** parameter, Directory Server logged all attributes into the error log file, including attributes containing sensitive information. A patch has been applied to filter out values of sensitive attributes. As a result, Directory Server no longer logs sensitive information. (BZ#1387772)

Group ACIs are now correctly evaluated

Previously, if the number of members in a group in an access control instruction (ACI) exceeded the size limit of the result of the query, Directory Server incorrectly denied access. To fix the problem, the server size limit is no longer applied to the ACI group evaluation, and queries now operate correctly. (BZ#1387022)

CHAPTER 7. FILE SYSTEMS

The autofs package now contains the `README.autofs-schema` file and an updated schema

The **samples/autofs.schema** distribution file was out of date and incorrect. As a consequence, it is possible that somebody is using an incorrect LDAP schema. However, a change of the schema in use cannot be enforced. With this update:

- The **README.autofs-schema** file has been added to describe the problem and recommend which schema to use, if possible.
- The schema included in the autofs package has been updated to **samples/autofs.schema.new**. (BZ#703846)

A stale `dentry` object is no longer left in the `dentry` cache after a rename operation

On an NFS file system, a stale **dentry** object was left in the **dentry** cache after a rename operation that replaced an existing object. As a consequence, if either the old or the new name contained 32 characters or more, the entry with the old name appeared accessible. The underlying source code has been modified to unhash the stale **dentry**. As a result, a rename operation no longer causes a stale **dentry** object to occur. (BZ#1080701)

`autofs` mounts no longer enter an infinite loop after reaching a shutdown state

If an **autofs** mount reached a shutdown state, and a mount request arrived and was processed before the mount-handling thread read the shutdown notification, the mount-handling thread exited without cleaning up the **autofs** mount. As a consequence, the main program never reached its exit condition and entered an infinite loop, as the autofs-managed mount was left mounted. To fix this bug, the exit condition check now takes place after each request is processed, and cleanup operations are now performed if an **autofs** mount has reached its shutdown state. As a result, the autofs daemon now exits as expected at shutdown. (BZ#1277033)

`automount` no longer needs to be restarted to access maps stored on the NIS server

Previously, the **autofs** utility did not wait for the NIS client service when starting. As a consequence, if the network map source was not available at program start, the master map could not be read, and the **automount** service had to be restarted to access maps stored on the NIS server. With this update, **autofs** waits until the master map is available to obtain a startup map. As a result, **automount** can access the map from the NIS domain, and **autofs** no longer needs to be restarted on every boot.

If the NIS maps are still not available after the configured wait time, the **autofs** configuration **master_wait** option might need to be increased. In the majority of cases, the wait time used by the package is sufficient. (BZ#1350786)

Setting the retry timeout can now prevent `autofs` from starting without mounts from SSSD

When starting the **autofs** utility, the **sss** map source was previously sometimes not ready to provide map information, but **sss** did not return an appropriate error to distinguish between a **map does not exist** and a **not available** condition. As a consequence, automounting did not work correctly, and **autofs** started without mounts from SSSD. To fix this bug, **autofs** retries asking SSSD for the master map when the **map does not exist** error occurs for a configurable amount of time. Now, you can set the retry timeout to a suitable value so that the master map is read and **autofs** starts as expected. (BZ#1384404)

CHAPTER 8. HARDWARE ENABLEMENT

Additional device IDs added to Intel NVMe driver

This patch adds the device IDs 0x0a53 and 0x0a54 to the Intel NVMe driver. This allows alignment of the I/O to the stripe size specified in the vendor specific controller **`Identify`** field, which can significantly improve performance. (BZ#1372088)

Fix for continuous probe of 82599ES when no SFP is installed

Using an Intel 82599ES card without SFP modules installed, could lead to excessive polling of the hardware and performance issues. The ixgbe driver has been updated to prevent this happening and to respond correctly should an SFP added at a later date. (BZ#1304849)

The `bnx2x` driver needs less time to recover after a parity event

Previously, the **bnx2x** driver waited for transmission completions when recovering from a parity event, which substantially increased the recovery time. With this update, **bnx2x** does not wait for transmission completion in the described circumstances. As a result, the recovery of **bnx2x** after a parity event now takes less time. (BZ#1300681)

Fix for inaccessible CIFS shares when using kerberos and multiuser

When reconnecting to a share that requires signed authentication, it is necessary to ensure that the client is able to sign requests if required. This patch ensures that signing will happen correctly. (BZ#1323053)

Firmware hangs with qlcnic driver

To prevent hangs, the driver has been updated to ensure that all writes are complete before the hardware fetches any pending transmission descriptors (BZ#1342659)

Updated microcode for AMD Processors

This release includes version(84F) of the AMD microcode, resolving issues with AMD Family Processors from 10h to 15h inclusively. (BZ#1322525)

Hang affecting raid1d when handling a mix of read and write errors

When hot removing or adding devices from a RAID1 array, the raid1d daemon could deadlock and become unresponsive in the freeze_array() function. This update includes bio_end_io_list writes in the nr_queued variable to prevent the hang from occurring. (BZ#1320595)

A race condition no longer occurs with IMSM RAID arrays running an `mdadm` reshape operation

With Intel Matrix Storage Manager (IMSM) RAID arrays running an **`mdadm`** reshape operation, a race condition could previously allow a second reshape to be launched on the same array before the first operation was completed, and the reshaping operation did not complete correctly. With this update, the race condition no longer occurs, and a second reshape operation cannot be started before the first operation is completed. (BZ#1340768)

Resolved kernel panic with Intel x520 FCoE hardware

Because FIP selection time was being reset for every FCoE controller advertisement received, FCF selection was being prevented.

As a result, when using FCoE as a root fs, the kernel would drop to a dracut shell with a panic message during boot.

This fix resets the FIP selection timer for the first authenticated controller advertisement only, allowing FCF selection to proceed. (BZ#1230719)

Resolution for large memory leak when using O_DIRECT I/O on an md device with DIX enabled

To prevent un-reclaimed memory, block IO integrity data is explicitly freed during the callback when bio is freed. (BZ#1268434)

SMM thermal interrupts are now handled properly

On some systems, enabling hardware p-states could cause a hang during boot due to improper handling of thermal interrupts from the SMM (System Management Mode). This patch allows the kernel to correctly handle the thermal interrupts. (BZ#1327715)

The kernel no longer panics after running the `ipmitool` command

In some cases, running the **`ipmitool`** command caused a kernel panic due to a race condition in the ipmi message handler. This update fixes the race condition, and the kernel panic no longer occurs in the described scenario. (BZ#1347189)

The operating system with SRP devices configured for auto startup on boot now boots correctly

Previously, a kernel crash on boot sometimes occurred when the operating system had the SCSI Remote Protocol (SRP) devices configured for auto startup on boot. This update fixes the srp_queuecommand() function, and the operating system now boots without problems in the described circumstances. (BZ#1321094)

The kernel no longer panics after running the `halt -p` command

When shutting down the system by running the **`halt -p`** command, a kernel panic occurred due to a conflict between the kernel offlining CPUs and the **sched** command, which used the sched group and the sched domain data without first checking the data. The underlying source code has been fixed by adding a check to avoid the conflict. As a result, the described scenario no longer produces a kernel panic. (BZ#1313035)

The ixgbe driver has been updated to the latest upstream version

This fix resoves a potential problem with stale pool mappings, which can result in MAC addresses being inadvertently cleared. (BZ#1346399)

Resolution for unavailability of shared IPMI on Intel 10G network cards

This updated driver ensures that the IPMI interface remains available when the network ports are powered down. (BZ#1316845)

Fix for kernel panic with HPSA drivers

Some configurations of HPSA devices caused hardware to be set offline incorrectly, when the driver should have waited for existing IO to complete, additional NMI code has been added to prevent this. (BZ#1083110)

Resolution for multiple ACPI errors on Intel CPUs

This fix addresses both ACPI namespace lookup errors and ACPI AE_AML_INFINITE_LOOP errors being repeatedly logged on some newer Intel CPUs. (BZ#1324697)

Resolution for VPD error messages in dmesg.

Some hardware incorrectly identifies itself as supporting VPD (Vital Product Data) information, when in fact it does not. This fix ensures the kernel now works around most cases of devices with poor or malformed VPD capabilities. (BZ#1289559)

Fixed performance problems when using Intel Xeon Coprocessor x100 product family and more than 255 CPUs

When Red Hat Enterprise Linux 6 was booted on a computer with Intel Xeon Coprocessor x100

Product Family with more than 255 CPUs, the kernel did not enumerate processor table entries in the correct order. Consequently, the system's performance was suboptimal. With this update, a patchset has been provided that ensures that the aforementioned entries are handled in the correct order, thus fixing this bug. (BZ#1247218)

Incorrect paths in scripts for IBM RSCT (Reliable Scalable Cluster Technology)

A bug in the update scripts for IBM RSCT caused the machine type and serial number to not be updated after migration. These paths have been corrected to enable migration and update. (BZ#1348279)

The weak-modules function now checks external symbols before failing a module as not compatible

When adding a module with a different version number to the current kernel, weak-modules checks for exported symbols in that kernel.

In some cases, the symbols required are provided by other modules. This fix ensures that the symbols provided by other modules are checked before marking the new module as incompatible. (BZ#1284935)

Fix for TRIM support being disabled on large RAID4/5/6 devices.

Due to an error calculating the limit of TRIM support, it was possible that larger RAID devices would disable TRIM.

With this update, max_discard_sectors and the volume stripe size are compared using the correct numerical units. TRIM support now works as expected. (BZ#1295969)

Resolution for VPD error messages in dmesg.

Some hardware incorrectly identifies itself as supporting VPD (Vital Product Data) information, when in fact it does not. This fix ensures the kernel now works around most cases of devices with poor or malformed VPD capabilities. (BZ#1324680)

Resolution for the alsaloop daemon consuming 100% CPU

Under some configurations, the alsaloop daemon could improperly use all CPU time for one or more CPUs. This fix resolves the issue by the addition of new error checking code. (BZ#1108292)

Resolution for IBM Power systems showing incorrect network link state

Using the **`ip addr`** command on IBM Power systems would incorrectly show some network hardware as having a link state of **`UNKNOWN`**. This has been fixed. (BZ#1089134)

Problem with obtaining Kerberos credentials after the session owner logs out no longer occurs

When mounting a Common Internet File System (CIFS) share using Kerberos authentication, the session key was previously revoked when the user owning the session logged out. Consequently, long running daemon processes started by the session were not able to use the request_key mechanism to obtain Kerberos credentials. With this update, the kernel packages have been fixed by creating a new keyring used to cache the CIFS Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO) upcalls. As a result, the session key no longer affects the SPNEGO keys, and the problem with obtaining Kerberos credentials no longer occurs. (BZ#1267754)

Resolution for dropped VLAN frames when using the e1000 driver

The bridge driver was failing to pass vlan registrations to the underlying physical devices. This resulted in frames for any vlans attached to a bridge were being lost.

The driver now calls vlan_rx_register() for any physical devices attached to a bridge when a vlan is modified, added, or deleted. This ensures VLANs and bridges work on RHEL 6. (BZ#1313663)

Resolution for kernel warning messages from vmxnet3 devices when softirq is disabled

Because of improper locking when softirq was disabled, warning messages and call traces were being captured in syslog.

The locking mechanism has been updated to not hold locks in this scenario, resolving the warnings. (BZ#1311537)

Fix for being unable to start system when using Trusted Boot (tboot) on a KVM machine

Because some hardware registers are not available to a KVM guest, tboot is unable to test for compatibility and returns a General Protection Fault (GPF). Although GPFs are reported, the user sees only a blank screen and no error message as the tboot process continues to retry access to the registers.

To resolve this, tboot now checks for Virtual Machine eXtensions (VMX) support before reading the feature control Model-Specific Register (MSR). (BZ#1323660)

Fix for failed install of tboot

The installation of tboot fails if the MANPATH environment variable is defined and has more than one directory

As MANPATH should not used in this manner, the installation directory has been changed to **`$(DISTDIR)/usr/share/man`** which resolves the issue. (BZ#1321857)

Fix for `Memory online failed` messages with VMware ESXi

When booting a guest on VMware ESXi hosts, an error is reported about enabling memory that has already been added.

The memory window (range) reported as an error is harmless, as both the E820 tables and the ACPI memory device are trying to add the ranges.

In normal usage these errors are redundant, the code has been changed to only report them if kernel verbosity is increased. (BZ#1255272)

Resolution for the `bnx2` driver using inappropriate spinlock functions

When netconsole was enabled, the bnx2 driver used spinlock_bh() and spinunlock_bh() in code that could be called in IRQ context. In this case, spinunlock_bh() incorrectly enabled interrupts. To fix this bug, the code has been altered to use spinlock_irqsave() and spinunlock_irqrestore() instead. (BZ#1291369)

Fix for Xen platforms being unable to select TSC as a clocksource

Additional code has been added to ensure that **`clocksource=tsc`** on the kernel command line will correctly select the kernel Time Stamp Counter (TSC) as the clock source. (BZ#1356231)

Additional code to support HP Pixart optical mouse

To prevent a problem, where some Hewlett Packard Pixart mice were unusable because of a constant connect and disconnect cycle, code has been backported from RHEL 7 and upstream. (BZ#1295575)

Fix for link flapping with `igbvf` driver and MSI-X interrupts

The **`igbvf`** driver was switching the carrier signal off every time an MSI-X interrupt was requested. This could be triggered by calling **`dhclient`** on the interface, or by configuring a multicast address.

As a result, the interface would be set down and up repeatedly during normal operation (link flapping), with log messages similar to:

```
kernel: igbvf 0000:00:05.0: Link is Up 1000 Mbps Full Duplex
```

This patch removes the **`carrier off`** operation when requesting an MSI-X interrupt, which removes the problem. (BZ#1304114)

Fix for system panic when booting with ixgbe driver configured with bonding and VLANs

With some configurations using device bonding and 802.1q VLANs, the ixgbe driver would cause a kernel panic during boot.

Code has been added to ensure that both the VLAN port number is correctly stored and the Answer to Reset (ATR) mechanism handles IPv6 extensions properly, preventing this behavior. (BZ#1339563)

Resolution for VPD error messages in dmesg.

Some hardware incorrectly identifies itself as supporting VPD (Vital Product Data) information, when in fact it does not. This fix ensures the kernel now works around most cases of devices with poor or malformed VPD capabilities. (BZ#1340999)

Include upstream code to fix a system crash due to an invalid pointer in CIFS

It was possible to crash a system in find_writable_file() because it called a version of use-after-free() that had errors.

This release incorporates upstream code to correct the issue. (BZ#1295008)

CHAPTER 9. INSTALLATION AND BOOTING

Installation on IBM System z with a network device in IPv6-only mode is now possible

Previously, when an IPv6-only network device was activated during initialization on IBM System z, a segmentation fault occurred and installation failed. The installer now handles these devices correctly, and no segmentation fault occurs, allowing installation to proceed. (BZ#1329171)

Adding a new EFI boot entry on a multipath device no longer fails

Previously, the **efibootmgr** utility was run multiple times for each device in multipath, which led to duplicate boot entries and possible failure. With this update, a new EFI boot entry is added only once for the whole multipath device. (BZ#1346725)

The size of the output buffer in `efibootmgr` has been increased

Previously, the output of the **efibootmgr** utility was sometimes truncated. The size of the output buffer has been increased, and now **efibootmgr** can show more detailed information. (BZ#1347283)

The `gateway` installation boot option now handles `IPv6` addresses

Previously, the **gateway** boot option in the installation system supported only **IPv4** addresses. With this update, it is possible to configure an **IPv6** gateway, using the **gateway** boot option, during installation. (BZ#1177984)

Thin provisioning can be successfully configured during installation

Due to changes in the way the LVM tools are packaged, the installation system raised an error while configuring thin provisioning. With this update, the installation environment is able to use these tools correctly, and thin provisioning can now be configured successfully. (BZ#1350296)

`ifdown` on a loopback device now works properly

In Red Hat Enterprise Linux version 6.7 and 6.8, executing the **ifdown** command on a local loopback device failed to remove the device. A patch has been applied, and the removal of an existing loopback device using **ifdown** now succeeds. (BZ#1311811)

Scripts in initscripts handle static IPv6 address assignment more robustly

Previously, scripts in the initscripts package sometimes failed to correctly assign static IPv6 addresses if a Router Advertisement (RA) was received during system initialization. This bug has been fixed, and now the statically assigned address is correctly applied in the described situation. (BZ#1086388)

Decompression of initrd larger than 32 MB no longer fails in GRUB on SGI UV100/UV1000 hardware

Previously, decompressing the **initrd** image in the **GRUB** boot loader could fail on SGI UV100 and UV1000 servers in cases where the image was loaded from a **TFTP** server and was larger than 32 MB. **GRUB** has been updated, and the **initrd** will now decompress successfully. (BZ#1124862)

The `ifup-aliases` script now sends gratuitous ARP updates when adding new IP addresses

When moving one or more IP aliases from one server to another, associated IP addresses may be unreachable for some time, depending on the Address Resolution Protocol (ARP) time-out value that is configured in the upstream router. This bug has been addressed in the initscripts package, and **ifup-aliases** now updates other systems on the network significantly faster in this situation. (BZ#1320366)

The initscripts package now handles LVM2 correctly

Previously, due to a bug, the initscripts package sometimes caused errors during boot on systems with using version 2 of the Logical Volume Manager device mapper (LVM2). With this update, this bug is fixed, and systems with LVM2 now boot correctly. (BZ#1331684)

The `netconsole` utility now launches correctly

Previously, if **`nameserver`** address lines were not present in the **`/etc/resolv.conf`** file, launching **`netconsole`** sometimes resulted in an error and **`netconsole`** did not start. The initscripts package has been updated, and **`netconsole`** now starts correctly in this situation. (BZ#1278519)

The `service network stop` command no longer attempts to stop services which are already stopped

Previously, when a tunnel interface was present, the **`service network stop`** command incorrectly attempted to stop services which had been stopped already, displaying an error message. With this update, this bug is fixed, and the **`service network stop`** command now only stops running services. (BZ#1315018)

The `dhclient` command no longer incorrectly uses `localhost` when `hostname` is not set

The **`dhclient`** command incorrectly sent **`localhost`** to the DHCP server as the host name when the **`hostname`** variable was not set. This has been fixed, and **`dhclient`** no longer sends an incorrect host name in these situations. (BZ#1350602)

CHAPTER 10. KERNEL

Reservation of memory for crashkernel no longer fails

Previously, the reservation of memory for crashkernel in some cases failed with the following error message:

```
Crashkernel reservation failed. Found area can not be reserved:
start=0x4000000, size=0x34000000.
```

This update fixes the step down mechanism so that the upper limit set in the **KEXEC_RESERVE_UPPER_LIMIT** parameter is not exceeded, which makes the reservation succeed. As a result, the memory reservation for crash kernel now proceeds as expected. (BZ#1349069)

The `mbind` call now allocates memory on the specified NUMA node

Previously, using the **mbind** call for allocation of memory on a Non-Uniform Memory Access (NUMA) node with particular number worked only for the very first invocation. On subsequent calls, the memory was always allocated on NUMA node 0. This update fixes the interaction of the mbind_range() function and the vma_adjust() function. As a result, **mbind** now allocates memory on the NUMA node with specified number in all cases. (BZ#1277241)

The system no longer hangs due to the `tasklist_lock` variable starvation

In a situation with a lot of concurrent processes taking the **tasklist_lock** variable for reading, the operating system sometimes became unresponsive when it was trying to take **tasklist_lock** for writing. This update fixes the underlying source code, so that a writer excludes the new readers to prevent the system hang. (BZ#1304864)

Intel Xeon v5 no longer causes GPU to hang

Previously, on GT3 and GT4 architectures, Intel Xeon v5 integrated graphics could experience problems with GPU lock-up, leading to GPU hang. This bug has been fixed. (BZ#1323945)

Kernel no longer panics when loading Intel Xeon v5 integrated graphic cards

When loading Intel Xeon v5 integrated graphic cards, a kernel panic sometimes occurred due to a race condition in the kernel firmware loader. This update adds a separate lock that is held throughout the life time of the firmware device, thus protecting the area where the device is registered. As a result, the kernel no longer panics in the described situation. (BZ#1309875)

NFS no longer uses FS-Cache when `-o fsc` is not set

Previously, when an NFS share was mounted, FS-Cache was always erroneously enabled even when the **-o fsc** option was not used in the mount command. Consequently, the **cachefilesd** service stored files on the NFS share, and other severe problems, such as the kernel becoming unresponsive or terminating unexpectedly, sometimes occurred. With this update, NFS no longer uses FS-Cache if **-o fsc** is not set. As a result, NFS now uses FS-Cache only when explicitly requested. Note that FS-Cache is provided as a Technology Preview in Red Hat Enterprise Linux 6. (BZ#1353844)

CHAPTER 11. NETWORKING

`ethtool -P` now returns correct output for virtual devices

The permanent address of virtual devices such as **`bridge`** or **`veth`** was incorrectly set to their link layer address, instead of being all zeroes. This update restores the original behavior for devices that do not have a permanent address so that the **`ethtool -P`** command now returns an all-zero address once again. (BZ#1318500)

Clients using IEEE802.1x-port-based authentication no longer lose connectivity

Previously, the **`wpa_supplicant`** service stopped responding to Extensible Authentication Protocol (EAP) Identity Request packets after **`wpa_supplicant`** was reloaded. As a consequence, clients using IEEE802.1x-port-based authentication lost connectivity until **`wpa_supplicant`** was restarted. With this update, the client is deauthenticated after reconfiguration. As a result, clients using IEEE802.1x-port-based authentication no longer lose connectivity in the described situation. (BZ#1359044)

UDP `iperf` over IPv6 `ESP` no longer causes kernel panic

Previously, a transform (XFRM) lookup could be performed on an already transformed destination cache entry (dst_entry) using UDP over IPv6 with a connected socket in conjunction with **`IPsec`** in Encapsulating Security Payload (ESP) transport mode. As a consequence, invalid IPv6 fragments were transmitted from the host or the kernel could terminate unexpectedly due to a socket buffer (SKB) underrun. With this update, XFRM lookup on already transformed **`dst_entry`** is not possible. Using UDP **`iperf`** over IPv6 **`ESP`** no longer causes invalid IPv6 fragments transmissions or a kernel panic. (BZ#1327680)

`tty_ldisc_flush()` no longer causes ISDN crashes

When providing certain services and using the Integrated Services Digital Network (ISDN), the system could terminate unexpectedly due to the call of the **`tty_ldisc_flush()`** function. The provided patch removes this call and the system no longer hangs in the described scenario. (BZ#1328115)

CHAPTER 12. SERVERS AND SERVICES

`httpd` no longer fails to start if there is a comment in the `Allow` directive

Comments had been disallowed at the end of the **Allow** directive in the Apache **mod_authz_host** module. Consequently, the **httpd** daemon failed to start with a syntax error if such a comment was included. With this update, it is possible to insert comments starting with the # (hash) sign, where # is interpreted as an end of line, in **Allow** or **Deny** directives. (BZ#1349546)

`db_verify` no longer causes `db4` to run out of free mutexes

Previously, the **db4** database did not correctly release all unused mutexes. When running the **db_verify** command on **db4** database files multiple times, **db4** quickly ran out of resources for mutex operations. Consequently, **db4** exited with the error message:

```
Unable to allocate memory for mutex; resize mutex region
```

leaving the database in an inconsistent state. This bug has been fixed, **db4** now correctly releases mutexes, and the described problem no longer occurs. (BZ#1272680)

The `OpenPegasus` CIM server is no longer automatically enabled

Previously, the **OpenPegasus** Common Information Model (CIM) server started automatically after installing the tog-pegasus package and rebooting the machine. Because **OpenPegasus** generates a self-signed SSL certificate, the automatic start caused **OpenPegasus** to fail certain security tests. This update modifies the **OpenPegasus** init script, so that the **OpenPegasus** CIM server has to be explicitly started by the user after tog-pegasus is installed. (BZ#1277655)

PAM authentication with `openwsman` now works correctly

Previously, the **openwsman** Pluggable Authentication Modules (PAM) configuration file contained non-existing PAM modules. As a consequence, **openwsman** PAM authentication in some cases failed. The configuration file has been modified to contain only existing and correct PAM modules. As a result, **openwsman** PAM authentication now works more reliably. (BZ#1152654)

The SFCB server now verifies that a WBEM port is available before attempting to use it

The Small Footprint Common Information Model (CIM) Broker (SFCB) uses the same Web-based Enterprise Management (WBEM) port as the OpenPegasus CIM server and other programs. Previously, if the SFCB server was started when the WBEM port was already used by another program, the SFCB server terminated unexpectedly with a segmentation fault. With this update, the SFCB server verifies upon the start whether the desired WBEM port is free to use and proceeds only if the port is free. If the WBEM port is not free, the SFCB server finishes gracefully and logs the reason for its termination. (BZ#907204)

`dstat` utility now displays data correctly when used with GPFS

Prior to this update, the **dstat** utility displayed no data when the **--gpfs-ops** option was used. This bug has been fixed, and **dstat** now displays data correctly when used with General Parallel File System (GPFS). (BZ#989779)

Evince now displays PostScript files again

Due to a bug, the **evince** document viewer failed to display content of PostScript files. A patch has been applied and **evince** now displays PostSrcipt files again. (BZ# 1410260)

The lwresd service no longer fails to reconnect to forwarders

Previously, when the **resolver-query-timeout** parameter in the **/etc/lwresd.conf** file was set too low, the **lwresd** service stopped querying unavailable DNS forwarder servers after a certain

amount of time. If one or more unavailable DNS servers came online, it was possible that **`lwresd`** never used a forwarder that was online. To fix the problem, the service now uses a default minimum value of 10 seconds, if the **`resolver-query-timeout`** parameter is set too low. As a result, all forwarders are queried until one of them responds. (BZ#1325081)

CHAPTER 13. STORAGE

`/proc/diskstats` no longer becomes corrupted

Partitions are protected by read-copy-update (RCU) for performance reasons and are not properly protected against race conditions in two circumstances:

- When partitions are modified while there are in-flight requests.
- When partitions overlap, which is possible for DOS extended and logical partitions.

As a consequence, some fields of the **`/proc/diskstats`** file could become corrupted. This update fixes the problem by caching the partition lookup in the request structure. As a result, **`/proc/diskstats`** no longer becomes corrupted in the described situations. (BZ#1273339)

`multipathd` no longer reports success after a failed device resizing

If the **`multipathd`** service failed to resize a device, **`multipathd`** did not reset the size back to the original value internally. As a consequence, on future attempts to resize a device, **`multipathd`** reported a success even when **`multipathd`** did not resize the device. If resizing fails, **`multipathd`** now reverts the size of the device back to the original value internally. As a result, **`multipathd`** now only reports success if a device is resized successfully. (BZ#1328077)

multipath no longer crashes due to `libdevmapper` version mismatches

Previously, the multipath code was not linking to the correct libraries during a part of compilation. As a consequence, if **`device-mapper-multipath`** was used with a newer version of the **`libdevmapper`** library than it was compiled with, multipath sometimes terminated unexpectedly. Now, multipath correctly links to the necessary libraries during compilation. As a result, multipath no longer crashes due to the library version mismatch. (BZ#1349376)

Failures on some devices no longer keep multipath from creating other devices

Previously, the **`multipath`** command sometimes failed to set up working devices because of failures on unrelated devices, as multipath quit early if it failed to get the information on any of the devices that multipath was trying to create. With this fix, multipath no longer quits early if it fails to get information on some of the devices and failures on some devices no longer keep multipath from creating others. (BZ#1343747)

multipath no longer modifies devices with a DM table type of multipath that were created by other programs

Previously, the multipath tools assumed that they were in charge of managing all Device Mapper (DM) devices with a multipath table. As a consequence, the **`multipathd`** service modified the tables of devices that were not created by the multipath tools. Now, multipath tools now only operate on devices with DM Universally Unique Identifiers (UUIDs) starting with **`mpath-`**, which is the UUID prefix that multipath uses on all the devices it creates. As a result, multipath no longer modifies devices with a DM table type of multipath that were created by other programs. (BZ#1364879)

Change now takes effect immediately after using `lvchange --zero n` against an active thin pool

Previously, when the **`lvchange --zero n`** command was used against an active thin pool, the change did not take effect until the next time the pool was deactivated. With this update, the change takes effect immediately. (BZ#1328245)

An incorrect exit status of `mdadm -IRs` no longer causes error messages at boot time

Previously, the load_container() function was incorrectly trying to load a container from the member array. As a consequence, the **`mdadm -IRs`** command incorrectly returned the **`1`** exit status, which led to error messages at boot time. The load_container() function has been modified to prevent loading a

container from a member array. As a result, error messages at boot time no longer occur. (BZ#1348925)

With IMSM, migrating two RAIDs in a container no longer causes both arrays to become degraded

The Intel Matrix Storage Manager (IMSM) does not allow a change to the RAID level of arrays in a container with two arrays. Previously, IMSM performed an array count check after disks were removed. As a consequence, changing, for example, RAID 1 to RAID 0 in a container with two RAIDs returned an error message, but a degraded RAID 1 was left. With this update, the array count check happens before the disk removal, and the described problem no longer occurs. (BZ#1413615)

The IMSM array is now correctly assembled and successfully started

Previously, the Intel Matrix Storage Manager (IMSM) `events` field was not set with a generation number. As a consequence, the mdadm utility sometimes re-assemblied a container with outdated metadata, and a failure occurred. With this update, the IMSM `events` field is correctly set with the generation number. As a result, the IMSM array is correctly assembled and successfully started. (BZ#1413937)

CHAPTER 14. SYSTEM AND SUBSCRIPTION MANAGEMENT

Subscription Manager no longer crashes when `nl_connect()` is unable to establish a connection

Previously, the **open_netlink()** function did not check the return code of the **nl_connect()** function, and the **get_etherinfo()** function did not check the return address of the **rtnl_link_alloc_cache()** function at the **python-ethtool** interface. Consequently, when **nl_connect()** was unable to establish a connection, Subscription Manager terminated unexpectedly with a segmentation fault. With this update, **open_netlink()** and **get_etherinfo()** check return values of **nl_connect()** and **rtnl_link_alloc_cache()** respectively, and Subscription Manager no longer crashes in the described scenario. (BZ#1032779)

ps no longer removes do_ and sys_ prefixes

Prior to this update, the **ps** command removed **do_** and **sys_** prefixes from the wait channel name when the **wchan** format option was used. As a conseqence, this caused ambiguities in some kernel function names. The bug has been fixed, and as a result the **ps** command no longer removes **do_** and **sys_** prefixes. (BZ#1322111)

CHAPTER 15. VIRTUALIZATION

`kdump` now works correctly with crashes caused in an interrupt context on Hyper-V

Previously, if a system crash was triggered in an interrupt context, the **kdump** utility was not able to capture the memory dump on guest virtual machines running on the Microsoft Hyper-V hypervisor. This update prevents the guest kernel from performing redundant tasks during the guest shutdown process caused by a crash. As a result, **kdump** can create a memory dump correctly in the described scenario. (BZ#1301903)

`virt-what` now detects IBM POWER LPARs

Previously, the **virt-what** utility did not recognize if the current system was running in an IBM POWER logical partition (LPAR). As a result, the **subscription-manager** utility displayed misleading licensing information for guest virtual machines running in an IBM POWER LPAR. With this update, **virt-what** was fixed and can now detect when the guest virtual machine is running in an IBM POWER LPAR. As a result, **subscription-manager** now provides correct subscription information in this scenario. (BZ#1312431)

Unaligned block I/O requests are now correctly detected

Previously, the length of all elements of the scatter/gather list of block I/O requests was not checked to be aligned according to the required alignment of a backing storage. As a result, in some configurations, notably when cache is set to none, guests issuing I/O requests with at least one element in the scatter/gather list that has an unaligned length would get I/O errors. With this update, QEMU now checks the length of all elements, in addition to the base address, and determines whether or not requests are correctly aligned. Consequently, unaligned requests are correctly detected, QEMU uses a correctly aligned bounce buffer for I/O, and the request can be completed successfully. (BZ#1321862)

All bridge network interfaces are now listed for new devices in virtual machines

Previously, not all bridge network interfaces were listed when a device was added to a virtual machine. With this update, all bridge network interfaces are properly detected using data provided by the **libvirt** service. As a result, all bridge network interfaces known to **libvirt** are listed when a device is added to a virtual machine. (BZ#1333290)

Network connectivity maintained when using rtl8139 device emulation

Previously, when using rtl8139 device emulation, the virtual device sometimes disabled packet reception. As a consequence, network connectivity was lost. With this update, the issue was resolved, and network connectivity is maintained. (BZ#1356924)

Quiescing disks after virtual disk migration no longer causes the guest to stop responding

When a high number of virtual disk migrations was active at the same time, the guest virtual machine in some cases became unresponsive, because the **QEMU** service was attempting to quiesce all disks on the guest. With this update, **QEMU** only quiesces the source disk whose migration is finishing, which prevents the problem from occurring. (BZ#876993)

`-S 0` for `qemu-img convert` now works correctly

Previously, specifying the **-S 0** option with **qemu-img convert** did not always explicitly write zeroes to all sparse areas in the output file. As a result, running **qemu-img convert -S 0** did not always result in a fully allocated output file. With this update, the command has been fixed, and running **qemu-img convert -S 0** always results in a fully allocated output file.

Note that the qcow2 format offers a different preallocation option which can convert images to fully allocated output files faster than using the **-S 0** option. However, with other formats, this can only be achieved with the **-S 0** option. (BZ#1297653)

Booting guests with multiple FC adapters on Hyper-V no longer causes a critical error

When booting a guest virtual machine that uses multiple virtual Fiber Channel (FC) adapters on the Microsoft Hyper-V hypervisor, the guest in some cases encountered a critical error. With this update, the **storsvc** driver correctly updates the reference count when loading an FC adapter, which prevents the described problem from occurring. (BZ#1316861)

Kernel memory dumps can now be captured when crashes occur on secondary CPUs

Previously, the wrong non-maskable interrupt (NMI) status was reported by Hyper-V hosts for Generation 2 guests. As a result, when a crash happened on a secondary CPU, it was not possible to capture a kernel memory dump. With this update, upstream patches were applied to avoid reading incorrect NMI statuses for Hyper-V Generation 2 guests. The **kdump** utility now works properly, regardless of the CPU that caused the crash. (BZ#1320215)

Unloading the hv_utils module no longer causes crashes

Previously, a bug in the hv_utils module was causing the kernel to crash when the module was removed shortly after loading. With this update, an upstream patch fixing the module unload path was applied to address the bug. As a result, unloading the hv_utils module no longer causes crashes. (BZ#1321259)

Localized `virt-manager` texts are all correctly translated

Previously, some of the texts in **virt-manager** were not translated. With this update, all texts in **virt-manager** are correctly translated. (BZ#1321729)

Unloading the storvsc module no longer causes kernel crashes for Hyper-V guests

Previously, a race condition in the kernel caused sporadic kernel crashes when the storvsc module was unloaded and then reloaded on Hyper-V guests. With this update, a fix was applied to the kernel SCSI code to eliminate the race condition. As a result, unloading the storvsc module no longer causes kernel crashes for Hyper-V guests. (BZ#1343105)

Creating bridge network interfaces from bond network interface in `virt-manager` now works as expected

When using the **virt-manager** interface to create a bridge network interface from existing bond network interfaces, the bond network configuration file became corrupted and configuration files for slave network interfaces on the guest virtual machine were deleted.

With this update, when creating a bridge network interface, the **python-virtinst** service creates an XML definition of the interface, with specification of interfaces that are included in the bridge. As a result, creating a bridge network interface in the described scenario is more reliable, and does not cause the virtual network to stop working. (BZ#1350683)

libguestfs now identifies operating systems in virtual machines where `/usr/` is not in the same partition as `/`

Previously, libguestfs did not recognize guest operating systems where the **/usr/** directory was not in the same partition as **/**. As a result, libguestfs tools such as **virt-sysprep** did not perform as expected. With this update, libguestfs recognizes guest operating systems even when the **/usr/** directory is not located in the same partition as **/**. Consequently, libguestfs tools such as **virt-sysprep** perform as expected. (BZ# 1388407)

Windows 8 virtual machines now shut down properly

Previously, when powering down a Windows 8 guest virtual machine with an 32-bit AMD and Intel architecture processor, the virtual machine became unresponsive on the Shutting Down screen. With the update of **virtio** and **qxl** drivers, this has been fixed. Consequently, the shutdown process

finishes as expected. (BZ#1271469)

PART II. TECHNOLOGY PREVIEWS

This chapter provides a list of all available Technology Previews in Red Hat Enterprise Linux 6.9.

Technology Preview features are currently not supported under Red Hat Enterprise Linux subscription services, may not be functionally complete, and are generally not suitable for production use. However, these features are included as a customer convenience and to provide the feature with wider exposure.

Customers may find these features useful in a non-production environment. Customers are also free to provide feedback and functionality suggestions for a Technology Preview feature before it becomes fully supported. Errata will be provided for high-severity security issues.

During the development of a Technology Preview feature, additional components may become available to the public for testing. It is the intention of Red Hat clustering to fully support Technology Preview features in a future release.

For information on Red Hat scope of support for Technology Preview features, see https://access.redhat.com/support/offerings/techpreview/.

CHAPTER 16. GENERAL UPDATES

A new module helping to upgrade the Tomcat server from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7

This update adds a new module to the preupgrade-assistant-el6toel7 package as a Technology Preview. The module helps to upgrade from Tomcat version 6.0.24 in Red Hat Enterprise Linux 6 to Tomcat version 7.0.x in Red Hat Enterprise Linux 7 and provides information about incompatibilities found in the system configuration. When using the module, which is recommended only on non-production machines, several automatic changes are made to the Tomcat configuration files during the postupgrade phase to prevent certain known issues. Note that in the supported scenario, users should remove the tomcat6 packages before upgrading.

A new module helping to upgrade Java OpenJDK 7 and Java OpenJDK 8 from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7

The preupgrade-assistant-el6toel7 package provides a new module that handles upgrades of Java OpenJDK 7 and Java OpenJDK 8 from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7. The module, available as a Technology Preview, informs users about possible requested actions and installs the expected equivalents of the original Java OpenJDK packages on the target system. Note that Java OpenJDK 6 and earlier versions are not handled by the in-place upgrade, but the module informs users about expected risks and required manual actions.

CHAPTER 17. AUTHENTICATION AND INTEROPERABILITY

Apache Modules for External Authentication

A set of Apache modules was added to Red Hat Enterprise Linux 6.6 as a Technology Preview. The `mod_authnz_pam`, `mod_intercept_form_submit`, and `mod_lookup_identity` Apache modules in the respective packages can be used by Web applications to achieve tighter interaction with external authentication and identity sources, such as Identity Management in Red Hat Enterprise Linux.

Simultaneous maintaining of TGTs for multiple KDCs

Kerberos version 1.10 added a new cache storage type, DIR:, which allows Kerberos to maintain Ticket Granting Tickets (TGTs) for multiple Key Distribution Centers (KDCs) simultaneously and auto-select between them when negotiating with Kerberized resources. Red Hat Enterprise Linux 6.4 and later includes SSSD enhanced to allow the users to select the DIR: cache for users that are logging in using SSSD. This feature is introduced as a Technology Preview.

Package: sssd-1.13.3

Cross-Forest Kerberos Trust Functionality in Identity Management

The Cross-Forest Kerberos Trust functionality provided by Identity Management (IdM) is included as a Technology Preview. This feature allows to create a trust relationship between an IdM and an Active Directory (AD) domain. This means that users from the AD domain can access resources and services from the IdM domain with their AD credentials. No data needs to be synchronized between the IdM and AD domain controllers; AD user are always authenticated against the AD domain controller and information about users is looked up without the need for synchronization.

This feature is provided by the optional ipa-server-trust-ad package. This package depends on features which are only available in `samba4`. Because samba4-* packages conflicts with the corresponding samba-* packages, all samba-* packages must be removed before ipa-server-trust-ad can be installed.

When the ipa-server-trust-ad package is installed, the `ipa-adtrust-install` utility must be run on all IdM servers and replicas to enable IdM to handle trusts. When this is done, a trust can be established from the command line using the `ipa trust-add` command or the IdM web UI. For more information, see the *Identity Management Guide* for Red Hat Enterprise Linux .

Note that Red Hat recommends to connect Red Hat Enterprise Linux 6 IdM clients to a Red Hat Enterprise Linux 7 IdM server for cross-forest trust capability. Trusts are fully supported on servers running Red Hat Enterprise Linux 7. Configuration with Red Hat Enterprise Linux 6 clients connected to a Red Hat Enterprise Linux 7 server for cross-forest trust is fully supported as well. In such setups, it is recommended to use the latest version of Red Hat Enterprise Linux 6 on the client side and the latest version of Red Hat Enterprise Linux 7 on the server side.

Packages: ipa-3.0.0 and samba-3.6.23

CHAPTER 18. COMPILER AND TOOLS

System Information Gatherer and Reporter (SIGAR)

The System Information Gatherer and Reporter (SIGAR) is a library and command-line tool for accessing operating system and hardware level information across multiple platforms and programming languages. In Red Hat Enterprise Linux 6.4 and later, SIGAR is considered a Technology Preview package.

Package: sigar-1.6.5-0.4.git58097d9

CHAPTER 19. FILE SYSTEMS

FS-Cache

FS-Cache in Red Hat Enterprise Linux 6 enables networked file systems (for example, NFS) to have a persistent cache of data on the client machine.

Package: cachefilesd-0.10.2-3

CHAPTER 20. KERNEL

Kernel Media support

The following features are presented as Technology Previews:

- The latest upstream video4linux
- Digital video broadcasting
- Primarily infrared remote control device support
- Various webcam support fixes and improvements

Package: kernel-2.6.32-696

Linux (NameSpace) Container [LXC]

Linux containers provide a flexible approach to application runtime containment on bare-metal systems without the need to fully virtualize the workload. Red Hat Enterprise Linux 6 provides application level containers to separate and control the application resource usage policies through cgroups and namespaces. This release includes basic management of container life-cycle by allowing creation, editing and deletion of containers using the **libvirt** API and the **virt-manager** GUI. Linux Containers are a Technology Preview.

Packages: libvirt-0.10.2-62, virt-manager-0.9.0-33

Diagnostic pulse for the fence_ipmilan agent, BZ# 655764

A diagnostic pulse can now be issued on the IPMI interface using the **`fence_ipmilan`** agent. This new Technology Preview is used to force a kernel dump of a host if the host is configured to do so. Note that this feature is not a substitute for the **`off`** operation in a production cluster.

Package: fence-agents-4.0.15-13

CHAPTER 21. NETWORKING

Mellanox SR-IOV Support

Single Root I/O Virtualization (SR-IOV) is now supported as a Technology Preview in the Mellanox **libmlx4** library and the following drivers:

- **mlx_core**
- **mlx4_ib** (InfiniBand protocol)
- **mlx_en** (Ethernet protocol)

Package: kernel-2.6.32-696

QFQ queuing discipline

In Red Hat Enterprise Linux 6, the **tc** utility has been updated to work with the Quick Fair Scheduler (QFQ) kernel features. Users can now take advantage of the new QFQ traffic queuing discipline from userspace. This feature is considered a Technology Preview.

Package: kernel-2.6.32-696

CHAPTER 22. SECURITY

TPM

TPM (Trusted Platform Module) hardware can create, store and use RSA keys securely (without ever being exposed in memory), verify a platform's software state using cryptographic hashes and more. The trousers and tpm-tools packages are considered a Technology Preview.

Packages: trousers-0.3.13-2, tpm-tools-1.3.4-2

CHAPTER 23. STORAGE

dm-era Device Mapper

The device-mapper-persistent-data package now provides tools to help use the new **dm-era** device mapper functionality released as a Technology Preview. The **dm-era** functionality keeps track of which blocks on a device were written within user-defined periods of time called an **era**. This functionality allows backup software to track changed blocks or restore the coherency of a cache after reverting changes.

DIF/DIX support

DIF/DIX, is a new addition to the SCSI Standard and a Technology Preview in Red Hat Enterprise Linux 6. 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 receive, and stores both the data and the checksum. Conversely, when a read occurs, the checksum can be checked by the storage device, and by the receiving HBA.

The DIF/DIX hardware checksum feature must only be used with applications that exclusively issue ***O_DIRECT*** I/O. These applications may use the raw block device, or the XFS file system in ***O_DIRECT*** mode. (XFS is the only file system that does not fall back to buffered I/O when doing certain allocation operations.) Only applications designed for use with ***O_DIRECT*** I/O and DIF/DIX hardware should enable this feature.

For more information, refer to section *Block Devices with DIF/DIX Enabled*in the Storage Administration Guide.

Package: kernel-2.6.32-696

LVM Application Programming Interface (API)

Red Hat Enterprise Linux 6 features the new LVM application programming interface (API) as a Technology Preview. This API is used to query and control certain aspects of LVM.

Package: lvm2-2.02.143-11

CHAPTER 24. VIRTUALIZATION

Performance monitoring in KVM guests

As a Technology Preview, **KVM** can virtualize a performance monitoring unit (vPMU) to allow virtual machines to use performance monitoring. Additionally it supports Intel's "architectural PMU" which can be live-migrated across different host CPU versions, using the **-cpu** host option.

The virtual performance monitoring feature allows virtual machine users to identify sources of performance problems in their guests, using their preferred pre-existing profiling tools that work on the host as well as the guest. Note that this is an addition to the existing ability to profile a KVM guest from the host.

Package:kernel-2.6.32-696

System monitoring using SNMP

As a Technology Preview, Red Hat Enterprise Linux 6 allows Simple Network Management Protocol (SNMP) to be used for system monitoring. This allows KVM hosts to send SNMP traps on events so that hypervisor events can be communicated to the user via standard SNMP protocol. In addition, SNMP is capable of performing basic virtual networking operations, such as starting and stopping the virtual domain.

Package:libvirt-snmp-0.0.2-5

Zero-copy compatibility for macvtap-vhost

The *macvtap-vhost* zero-copy capability is available on Red Hat Enterprise Linux 6 as a Technology Preview. This feature allows running networking work loads in very high wire speeds but with low CPU resource consumption, and it does not limit other features such as memory overcommit and guest migration, which is not the case when using PCI device assignment to achieve the wire speed. Note that this feature is disabled by default.

Package:qemu-kvm-0.12.1.2-2.499

vCPU hot unplug

Although hot-plugging a virtual CPU (vCPU) is a supported operation, hot-unplugging a vCPU remains a Technology Preview in Red Hat Enterprise Linux 6, and is strongly recommended not to be used in high-value deployments.

Package:libvirt-0.10.2-62

CHAPTER 25. DEPRECATED FUNCTIONALITY

This chapter provides an overview of functionality that has been deprecated, or in some cases removed, in all minor releases up to Red Hat Enterprise Linux 6.9.

Deprecated functionality continues to be supported until the end of life of Red Hat Enterprise Linux 6. Deprecated functionality will likely not be supported in future major releases of this product and is not recommended for new deployments. For the most recent list of deprecated functionality within a particular major release, refer to the latest version of release documentation.

Deprecated *hardware* components are not recommended for new deployments on the current or future major releases. Hardware driver updates are limited to security and critical fixes only. Red Hat recommends replacing this hardware as soon as reasonably feasible.

A *package* can be deprecated and not recommended for further use. Under certain circumstances, a package can be removed from a product.Product documentation then identifies more recent packages that offer functionality similar, identical, or more advanced to the one deprecated, and provides further recommendations.

Deprecated Insecure Algorithms and Protocols

Algorithms that provide cryptographic hashes and encryption as well as cryptographic protocols have a lifetime after which they are considered either too risky to use or plain insecure. See the Deprecation of Insecure Algorithms and Protocols in RHEL 6.9 article on the Red Hat Customer Portal for more information.

MD5, MD4, and SHA0 can no longer be used as signing algorithms in OpenSSL

With this update, support for verification of MD5, MD4, and SHA0 signatures in certificates, Certificate Revocation Lists (CRL) and message signatures are removed.

The system administrator can enable MD5, MD4, or SHA0 support by modifying the **LegacySigningMDs** option in the **etc/pki/tls/legacy-settings** policy configuration file, for example:

```
echo 'LegacySigningMDs algorithm' >> /etc/pki/tls/legacy-settings
```

To add more than one legacy algorithm, use a comma or any whitespace character except a new line. See the **README.legacy-settings** in the **OpenSSL** package for more information.

You can also enable MD5 verification by setting the **OPENSSL_ENABLE_MD5_VERIFY** environment variable.

OpenSSL clients no longer allow connections to servers with DH shorter than 1024 bits

This change prevents **OpenSSL** clients from connecting to servers with Diffie-Hellman (DH) parameters shorter than 1024 bits. This ensures that allowed clients using **OpenSSL** are not vulnerable to attacks such as the LOGJAM attack.

The system administrator can enable shorter DH parameter support by modifying the **MinimumDHBits** option in the **/etc/pki/tls/legacy-settings**, for example:

```
echo 'MinimumDHBits 768' > /etc/pki/tls/legacy-settings
```

This option can also be used to raise the minimum if required by the system administrator.

EXPORT cipher suites in OpenSSL are deprecated

This change removes support for EXPORT cipher suites in the **OpenSSL** toolkit. Disabling these weak cipher suites prevents attacks such as the FREAK attack. EXPORT cipher suites are not required in any **TLS** protocol configuration.

GnuTLS clients no longer allow connections to servers with DH shorter than 1024 bits

This change prevents GNU Transport Layer Security (GnuTLS) clients from connecting to servers with Diffie-Hellman (DH) parameters shorter than 1024 bits. This ensures that allowed clients using **GnuTLS** are not vulnerable to attacks such as the LOGJAM attack.

The system administrator can enable shorter DH parameter support by modifying the **MinimumDHBits** option in the **/etc/pki/tls/legacy-settings**, for example:

```
echo 'MinimumDHBits 768' > /etc/pki/tls/legacy-settings
```

This option can also be used to raise the minimum if required by the system administrator.

EXPORT cipher suites in GnuTLS are deprecated

This change removes support for EXPORT cipher suites in the GNU Transport Layer Security (GnuTLS) library. Disabling these weak cipher suites prevents attacks such as the FREAK attack. EXPORT cipher suites are not required in any **TLS** protocol configuration.

The **GnuTLS** EXPORT cipher suite priority string remains, but as an alias for the NORMAL priority string.

MD5 can no longer be used as a signing algorithm in NSS

This change prevents the Network Security Services (NSS) library from using MD5 as the signing algorithm in **TLS**. This change ensures that programs using **NSS** are not vulnerable to attacks such as the SLOTH attack.

The system administrator can enable MD5 support by modifying the **/etc/pki/nss-legacy/nss-rhel6.config** policy configuration file to:

```
library=
name=Policy
NSS=flags=policyOnly,moduleDB
config="allow=MD5"

```

Note that an empty line is required at the end of the file.

NSS clients using TLS no longer allow connections to servers with DH shorter than 1024 bits

This change prevents Network Security Services (NSS) clients from connecting to servers with Diffie-Hellman (DH) parameters shorter than 1024 bits. This ensures that allowed clients using **NSS** are not vulnerable to attacks such as the LOGJAM attack.

The system administrator can enable shorter DH parameter support by modifying the **/etc/pki/nss-legacy/nss-rhel6.config** policy configuration file to:

```
library=
name=Policy
NSS=flags=policyOnly,moduleDB

```

```
config="allow=DH-MIN=767:DSA-MIN=767:RSA-MIN=767"
```

Note that an empty line is required at the end of the file.

EXPORT cipher suites in NSS are deprecated

This change removes support for EXPORT cipher suites in the Network Security Services (NSS) library. Disabling these weak cipher suites prevents attacks such as the FREAK attack. EXPORT cipher suites are not required in any **TLS** protocol configuration.

Deprecated algorithms in OpenSSH: RC4, hmac-md5, and hmac-md5-96

With this update, the **arcfour256**, **arcfour128**, **arcfour** ciphers and the **hmac-md5**, **hmac-md5-96** Message Authentication Code (MAC) algorithms are deprecated. Note that this change does not affect any existing server configuration.

The system administrator can enable these deprecated algorithms by editing the **ssh_config** file, for example:

```
Host legacy-system.example.com
  Ciphers arcfour
  MACs hmac-md5
```

To completely restore all the deprecated algorithms, add the following snippet to the **/etc/ssh/ssh_config** file:

```
Ciphers aes128-ctr,aes192-ctr,aes256-ctr,arcfour256,arcfour128,aes128-
cbc,3des-cbc,blowfish-cbc,cast128-cbc,aes192-cbc,aes256-
cbc,arcfour,rijndael-cbc@lysator.liu.se
MACs hmac-md5,hmac-sha1,umac-64@openssh.com,hmac-sha2-256,hmac-sha2-
512,hmac-ripemd160,hmac-ripemd160@openssh.com,hmac-sha1-96,hmac-md5-96
```

GnuTLS no longer provides cryptographic back-end replacement APIs

The functions implementing cryptographic back-end replacement are considered obsolete and act as no-operation functions now. The following functions exported in the **gnutls/crypto.h** file are affected:

- gnutls_crypto_single_cipher_register2
- gnutls_crypto_single_mac_register2
- gnutls_crypto_single_digest_register2
- gnutls_crypto_cipher_register2
- gnutls_crypto_mac_register2
- gnutls_crypto_digest_register2
- gnutls_crypto_rnd_register2
- gnutls_crypto_pk_register2
- gnutls_crypto_bigint_register2

Deprecated Drivers

Deprecated device drivers

- 3w-9xxx
- 3w-sas
- 3w-xxxx
- aic7xxx
- i2o
- ips
- megaraid_mbox
- mptbase
- mptctl
- mptfc
- mptlan
- mptsas
- mptscsih
- mptspi
- sym53c8xx
- qla3xxx

The following controllers from the **megaraid_sas** driver have been deprecated:

- Dell PERC5, PCI ID 0x15
- SAS1078R, PCI ID 0x60
- SAS1078DE, PCI ID 0x7C
- SAS1064R, PCI ID 0x411
- VERDE_ZCR, PCI ID 0x413
- SAS1078GEN2, PCI ID 0x78

The following controllers from the **be2iscsi** driver have been deprecated:

- BE_DEVICE_ID1, PCI ID 0x212
- OC_DEVICE_ID1, PCI ID 0x702
- OC_DEVICE_ID2, PCI ID 0x703

Note that other controllers from the mentioned drivers that are not listed here remain unchanged.

Other Deprecated Components

`cluster`, `luci` components

The `fence_sanlock` agent and `checkquorum.wdmd`, introduced in Red Hat Enterprise Linux 6.4 as a Technology Preview and providing mechanisms to trigger the recovery of a node using a hardware watchdog device, are considered deprecated.

openswan component

The openswan packages have been deprecated, and libreswan packages have been introduced as a direct replacement for openswan to provide the VPN endpoint solution. openswan is replaced by libreswan during the system upgrade.

`seabios` component

Native KVM support for the S3 (suspend to RAM) and S4 (suspend to disk) power management states has been discontinued. This feature was previously available as a Technology Preview.

The `zerombr yes` Kickstart command is deprecated

In some earlier versions of Red Hat Enterprise Linux, the `zerombr yes` command was used to initialize any invalid partition tables during a Kickstart installation. This was inconsistent with the rest of the Kickstart commands due to requiring two words while all other commands require one. Starting with Red Hat Enterprise Linux 6.7, specifying only `zerombr` in your Kickstart file is sufficient, and the old two-word form is deprecated.

Btrfs file system

B-tree file system (Btrfs) is considered deprecated for Red Hat Enterprise Linux 6. Btrfs was previously provided as a Technology Preview, available on AMD64 and Intel 64 architectures.

eCryptfs file system

eCryptfs file system, which was previously available as a Technology Preview, is considered deprecated for Red Hat Enterprise Linux 6.

`mingw` component

Following the deprecation of Matahari packages in Red Hat Enterprise Linux 6.3, at which time the mingw packages were noted as deprecated, and the subsequent removal of Matahari packages from Red Hat Enterprise Linux 6.4, the mingw packages were removed from Red Hat Enterprise Linux 6.6 and later.

The mingw packages are no longer shipped in Red Hat Enterprise Linux 6 minor releases, nor will they receive security-related updates. Consequently, users are advised to uninstall any earlier releases of the mingw packages from their Red Hat Enterprise Linux 6 systems.

`virtio-win` component, BZ#1001981

The VirtIO SCSI driver is no longer supported on Microsoft Windows Server 2003 platform.

`fence-agents` component

Prior to Red Hat Enterprise Linux 6.5 release, the Red Hat Enterprise Linux High Availability Add-On was considered fully supported on certain VMware ESXi/vCenter versions in combination with

the fence_scsi fence agent. Due to limitations in these VMware platforms in the area of SCSI-3 persistent reservations, the **fence_scsi** fencing agent is no longer supported on any version of the Red Hat Enterprise Linux High Availability Add-On in VMware virtual machines, except when using iSCSI-based storage. See the Virtualization Support Matrix for High Availability for full details on supported combinations: https://access.redhat.com/site/articles/29440.

Users using **fence_scsi** on an affected combination can contact Red Hat Global Support Services for assistance in evaluating alternative configurations or for additional information.

systemtap component

The systemtap-grapher package has been removed from Red Hat Enterprise Linux 6. For more information, see https://access.redhat.com/solutions/757983.

matahari component

The **Matahari** agent framework (matahari-*) packages have been removed from Red Hat Enterprise Linux 6. Focus for remote systems management has shifted towards the use of the CIM infrastructure. This infrastructure relies on an already existing standard which provides a greater degree of interoperability for all users.

distribution component

The following packages have been deprecated and are subjected to removal in a future release of Red Hat Enterprise Linux 6. These packages will not be updated in the Red Hat Enterprise Linux 6 repositories and customers who do not use the MRG-Messaging product are advised to uninstall them from their system.

- python-qmf
- python-qpid
- qpid-cpp
- qpid-qmf
- qpid-tests
- qpid-tools
- ruby-qpid
- saslwrapper

Red Hat MRG-Messaging customers will continue to receive updated functionality as part of their regular updates to the product.

fence-virt component

The **libvirt-qpid** is no longer part of the fence-virt package.

openscap component

The openscap-perl subpackage has been removed from openscap.

APPENDIX A. LIST OF BUGZILLAS BY COMPONENT

Table A.1. List of Bugzillas by Component

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
389-ds-base	BZ#1330758	BZ#1404443	BZ#1266920, BZ#1327065, BZ#1352109, BZ#1369572, BZ#1370145, BZ#1371678, BZ#1371706, BZ#1382386, BZ#1387022, BZ#1387772, BZ#1399600, BZ#1402012, BZ#1403754, BZ#1406835, BZ#1410645
NetworkManager	BZ#1308730		
alsa-utils			BZ#1108292
anaconda		BZ#914637, BZ#1014425, BZ#1253223, BZ#1416653	BZ#1177984, BZ#1329171, BZ#1346725, BZ#1350296
audit	BZ#1369249		
autofs			BZ#703846, BZ#1277033, BZ#1350786, BZ#1384404
bind			BZ#1325081
ca-certificates	BZ#1368996		
cloud-init	BZ#1421281		
clufter	BZ#1367536		
control-center			BZ#1217790

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
cpuid	BZ#1316998		
cups		BZ#1099617, BZ#1268131	
db4			BZ#1272680
device-mapper-multipath	BZ#1305589, BZ#1310320, BZ#1333334, BZ#1355669, BZ#1377532		BZ#1328077, BZ#1343747, BZ#1349376, BZ#1364879
dhcp	BZ#1321945	BZ#1297445	
distribution	BZ#1339222		
dstat			BZ#989779
efibootmgr			BZ#1347283
gdb			BZ#1316539
gdm			BZ#1083680, BZ#1337067
ghostscript		BZ#1411843	BZ#1410260
git			BZ#874659
glibc	BZ#1101858		BZ#1012343, BZ#1223095, BZ#1270950, BZ#1331304, BZ#1338673, BZ#1373646
gnome-session			BZ#1320245
gnome-settings-daemon			BZ#966658
grub			BZ#1124862
httpd			BZ#1349546

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
initscripts	BZ#1157856	BZ#1090559	BZ#1086388, BZ#1278519, BZ#1311811, BZ#1315018, BZ#1320366, BZ#1331684, BZ#1350602
ipa	BZ#1367026	BZ#1399058	

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
kernel	BZ#1167938, BZ#1306457, BZ#1306469, BZ#1343743, BZ#1347825, BZ#1349112, BZ#1352824, BZ#1365049, BZ#1392941	BZ#822725, BZ#1012684, BZ#1111683, BZ#1121888, BZ#1224673, BZ#1288597, BZ#1315832, BZ#1396336	BZ#1080701, BZ#1083110, BZ#1089134, BZ#1230719, BZ#1247218, BZ#1255272, BZ#1267754, BZ#1268434, BZ#1273339, BZ#1277241, BZ#1289559, BZ#1291369, BZ#1293709, BZ#1295008, BZ#1295575, BZ#1295969, BZ#1300681, BZ#1301903, BZ#1304114, BZ#1304849, BZ#1304864, BZ#1309875, BZ#1311537, BZ#1313035, BZ#1313663, BZ#1316845, BZ#1316861, BZ#1318500, BZ#1320215, BZ#1320595, BZ#1321094, BZ#1321259, BZ#1323053, BZ#1323945, BZ#1324680, BZ#1324697, BZ#1327680, BZ#1327715, BZ#1328115, BZ#1339563, BZ#1340999, BZ#1342659, BZ#1343105, BZ#1346399, BZ#1347189, BZ#1349069, BZ#1353844, BZ#1356231, BZ#1372088

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
krb5	BZ#1351284		
libcacard		BZ#1331471	
libguestfs			BZ#1388407
libvirt	BZ#1333415		
luci	BZ#885028, BZ#1173942		
lvm2			BZ#1328245
mdadm			BZ#1340768, BZ#1348925, BZ#1413615, BZ#1413937
microcode_ctl			BZ#1322525
module-init-tools			BZ#1284935
mutt			BZ#1196787
nss-pam-ldapd		BZ#1401632	
openldap			BZ#1249092
openscap	BZ#1364207		
openwsman			BZ#1152654
other	BZ#1318326	BZ#1336548	
pacemaker	BZ#1253325		BZ#1322595
pam	BZ#1404832		
perl			BZ#1364206
perl-Frontier-RPC			BZ#832390
perl-IO-Socket-SSL	BZ#1331037		

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
perl-Net-SSLeay	BZ#1325407		
perl-libwww-perl			BZ#1400632
pki-core			BZ#1403943
postfix	BZ#1287192		
ppc64-diag			BZ#1348279
preupgrade-assistant	BZ#1427713		
preupgrade-assistant-el6toel7	BZ#1392018, BZ#1402478, BZ#1406464	BZ#1366671, BZ#1388967	
procps			BZ#1322111
psacct			BZ#1182317
python-ethtool			BZ#1032779
python-virtinst			BZ#1350683
qemu-kvm		BZ#1063124, BZ#1198956, BZ#1209362, BZ#1271469, BZ#1346153, BZ#1371765	BZ#876993, BZ#1297653, BZ#1321862, BZ#1356924
radvd		BZ#1058698	
rear		BZ#1313417, BZ#1320551	
resource-agents	BZ#1336846		
rgmanager			BZ#1084053, BZ#1228170, BZ#1342825, BZ#1414139
rsyslog7	BZ#1323199		

Component	Release Notes		Technical Notes
	New Features	Known Issues	Notable Bug Fixes
ruby			BZ#1331086
sblim-sfcb			BZ#907204
scap-security-guide	BZ#1311491		
sssd	BZ#1324428, BZ#1329378		BZ#1293168, BZ#1321884, BZ#1327272, BZ#1335400, BZ#1336453, BZ#1340176, BZ#1367435, BZ#1374813
tboot			BZ#1321857, BZ#1323660
tcsh			BZ#885901, BZ#1334751, BZ#1338986
tog-pegasus			BZ#1277655
virt-manager			BZ#1321729, BZ#1333290
virt-what			BZ#1312431
virtio-win	BZ#1303906		
vsftpd	BZ#1350724		
wpa_supplicant			BZ#1359044
xorg-x11-drv-vmmouse		BZ#1322712	
xorg-x11-drv-vmware		BZ#1320480	
xorg-x11-server		BZ#1076595	

APPENDIX B. REVISION HISTORY

Revision 0.1-5 **Fri May 12 2017** **Lenka Špačková**

Moved the `fence_sanlock` agent and `checkquorum.wdmd` from Technology Previews to Deprecated Functionality.

Revision 0.1-3 **Thu Apr 27 2017** **Lenka Špačková**

Added the deprecated `zerombr yes` Kickstart command to Deprecated Functionality.

Revision 0.1-2 **Fri Mar 31 2017** **Lenka Špačková**

Added a bug fix to Virtualization.

Revision 0.1-1 **Tue Mar 28 2017** **Lenka Špačková**

Minor edits in accordance with the updated Red Hat Enterprise Linux 6.9 Release Notes.

Revision 0.0-9 **Tue Mar 21 2017** **Lenka Špačková**

Release of the Red Hat Enterprise Linux 6.9 Technical Notes.

Revision 0.0-5 **Thu Jan 05 2017** **Lenka Špačková**

Release of the Red Hat Enterprise Linux 6.9 Beta Technical Notes.