



Red Hat Enterprise Linux 6 6.9 Release Notes

Release Notes for Red Hat Enterprise Linux 6.9
Edition 9

Red Hat Customer Content Services

Red Hat Enterprise Linux 6 6.9 Release Notes

Release Notes for Red Hat Enterprise Linux 6.9 Edition 9

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

Legal Notice

Copyright © 2017 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

Preface	4
Chapter 1. Overview	5
Production Phase	5
In-place Upgrade	5
Security	5
Red Hat Insights	5
Red Hat Access Labs	5
Part I. New Features	7
Chapter 2. General Updates	8
In-place upgrade from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7	8
preupgrade-assistant rebased to version 2.3.3	8
Preupgrade Assistant enables blacklisting to improve performance	8
Key file names unified in Preupgrade Assistant modules	9
A new RHDS module to check a possibility of an in-place upgrade of an RHDS system	9
cloud-init moved to the Base channel	9
Chapter 3. Authentication and Interoperability	10
SSSD now enables the administrator to select which domains from the AD forest can be contacted	10
SSSD now enables selecting a list of PAM services that will not receive any environmental variables from pam_sss	10
IdM servers can now be configured to require TLS 1.2 or better	10
pam_faillock can be now configured with unlock_time=never	10
The libkadm5* libraries have been moved to the libkadm5 package	10
Chapter 4. Clustering	12
Support added for Oracle 11g in Oracle and OrLsnr Pacemaker resource agents	12
Pacemaker now supports alert agents	12
clutter is now fully supported	12
clutter rebased to version 0.59.8	12
luci interface allows administrators to verify authenticity of remote machines	12
luci now lists explicit configured actions for individual resources	13
Chapter 5. Compiler and Tools	14
Support for the el_GR@euro, ur_IN, and wal_ET locales has been added	14
The Net:SSLeay Perl module now supports restricting of TLS version	14
The IO::Socket::SSL Perl module now supports restricting of TLS version	14
ca-certificates rebased to version 2.10	14
Chapter 6. Directory Server in Red Hat Enterprise Linux	15
Directory Server now supports enabling and disabling specific TLS versions	15
Chapter 7. Hardware Enablement	16
cpuid is now available	16
Support for RealTek RTS5250S SD4.0 Controllers	16
Chapter 8. Installation and Booting	17
The NO_DHCP_HOSTNAME option has been added	17
Chapter 9. Kernel	18
Chelsio firmware updated to version 1.15.37.0	18
The bnx2_en driver updated to the latest upstream version	18
The ahci driver supports Marwell 88SE9230	18

.....	..
Chapter 10. Networking	19
NetworkManager now supports manual DNS configuration with dns=none	19
Chapter 11. Security	20
TLS 1.2 support added to all system components	20
OpenSCAP 1.2.13 is NIST certified	20
vsftpd now uses TLS 1.2 by default	20
auditd now supports incremental_async	20
scap-security-guide now supports ComputeNode	20
rsyslog7 now enables TLS 1.2	20
Chapter 12. Servers and Services	21
A DHCP client hook example added for DDNS for Microsoft Azure cloud	21
postfix now supports user-controlled configuration of TLS	21
Chapter 13. Storage	22
The smartPQI (smartpqi) driver is now available	22
Update of mpt3sas	22
Update of megaraid_sas	22
A new default configuration for Huawei XSG1 arrays has been added for device-mapper-multipath	
The disable_changed_wwids multipath.conf option is now available in multipath to avoid data corruption	22
device-mapper-multipath now supports the max_sectors_kb configuration parameter	23
The skip_kpartx multipath.conf option to allow skipping kpartx partition creation has been added	
Users are now warned if they create multipath devices while multipathd is not running	23 23
Chapter 14. Virtualization	24
Configuration options can be used to exclude weak ciphers	24
Improved Hyper-V storage driver performance	24
Hyper-V clock source changed to use the TSC page	24
Setting the account password is now possible for any guest user	24
Chapter 15. Red Hat Software Collections	25
Part II. Known Issues	26
Chapter 16. General Updates	27
The default value of first_valid_uid in Dovecot has changed in Red Hat Enterprise Linux 7	27
Incorrect information about the expected default settings of services in Red Hat Enterprise Linux 7	27
Chapter 17. Authentication and Interoperability	28
SSSD fails to manage sudo rules from the IdM LDAP tree	28
winbindd crashes when installing a new AD trust	28
nslcd fails to resolve user or group identities when it is started before the network connection is fully up	28
Chapter 18. Desktop	29
The vmware driver does not support multiple displays	29
Incorrect mouse pointer movement after screen rotation inside a virtual machine in VMWare 11 or VMWare 12	29
Using Radeon or Nouveau can cause incorrectly rendered graphics	29
Chapter 19. Directory Server in Red Hat Enterprise Linux	30
IdM schema replications from Red Hat Enterprise Linux 7 to 6.9 fail	30

Chapter 20. Installation and Booting	31
The installer displays the number of multipath devices, and number of multipath devices selected, incorrectly	31
The installer displays the amount of disk space within multipath devices incorrectly	31
The device.map configuration file generated by Anaconda is sometimes incorrect	31
The ifup script incorrectly replaces manually-defined default routes	31
Upgrading Red Hat Enterprise Linux 6 on UEFI systems clears the boot loader password	31
Chapter 21. Kernel	32
Certain NIC firmware can become unresponsive with the bnx2x driver	32
e1000e cards might not get an IPv4 address	32
The ecb kernel module fails when dracut is not upgraded	32
Chapter 22. Networking	33
The radvd occasionally terminates unexpectedly due to a race condition	33
Chapter 23. Servers and Services	34
Printing a PDF file upside down with cups is currently impossible	34
Printing PDF files using the fit-to-page and fitplot options does not work on printers with hardware margins	34
DHCP client sends unicast requests through the incorrect interface	34
A *.dsc file converted from a *.pdf file by the pdf2dsc script cannot be opened in Evince	34
Chapter 24. Storage	35
mkfs might become temporarily unresponsive on the 32-bit AMD and Intel architecture	35
Chapter 25. System and Subscription Management	36
ReaR works only on the eth0 interface	36
ReaR creates two ISO images instead of one	36
Chapter 26. Virtualization	37
Coolkey does not load on Windows 7 guests	37
Disabling vCPUs on Hyper-V guests fails	37
Hot plugging hard disks as a batch on the VMware ESXi hypervisor does not work reliably	37
Guests cannot access floppy disks larger than 1.44 MB	37
Hyper-V guest integration services stop working after they are disabled and re-enabled	37
Booting virtual machines with the fsgsbase and smep flags on older host CPUs fails	37
Guests with recent Windows systems in some cases fail to boot if hv_relaxed is used	38
Limited CPU support for Windows 10 and Windows Server 2016 guests	38
Windows 8 Virtual Machine Now Shuts Down Properly	38
Network connectivity not restored when vnic is enabled	38
Appendix A. Component Versions	39
Appendix B. Revision History	40

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 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. The [Technical Notes](#) document provides a list of notable bug fixes, all currently available Technology Previews, deprecated functionality, and other information.

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. Overview

Production Phase

Red Hat Enterprise Linux 6 is now in Production Phase 2, and Red Hat Enterprise Linux 6.9 therefore provides a stable release focused on bug fixes. Red Hat Enterprise Linux 6 enters Production Phase 3 on May 10, 2017. Subsequent updates will be limited to qualified critical security fixes and business-impacting urgent issues. Please refer to [Red Hat Enterprise Linux Life Cycle](#) for more information.

In-place Upgrade

As Red Hat Enterprise Linux subscriptions are not tied to a particular release, existing customers can update their Red Hat Enterprise Linux 6 infrastructure to Red Hat Enterprise Linux 7 at any time, free of charge, to take advantage of recent upstream innovations. To simplify the upgrade to Red Hat Enterprise Linux 7, Red Hat provides the **Preupgrade Assistant** and **Red Hat Upgrade Tool**. For more information, see [Chapter 2, General Updates](#).

Security

- With the addition of TLS protocol version 1.2 support to the **GnuTLS** component, Red Hat Enterprise Linux 6 offers complete support for TLS 1.2 in the provided security libraries. TLS 1.2 is recommended by modern security standards such as PCI-DSS 3.1. For more information, see [Chapter 11, Security](#).
- **OpenSCAP** 1.2.13 has been certified by the National Institute of Standards and Technology's (NIST) Security Content Automation Protocol (SCAP) 1.2 in the Authenticated Configuration Scanner category with the Common Vulnerabilities and Exposure (CVE) option. For details, see [Chapter 11, Security](#).
- Cryptographic protocols and algorithms that are considered insecure, such as MD5, SHA0, RC4, or DH shorter than 1024 bits, have been deprecated. In addition, support for EXPORT cipher suites has been removed. For details, see the [Red Hat Enterprise Linux 6.9 Technical Notes](#).

Red Hat Insights

Since Red Hat Enterprise Linux 6.7, the *Red Hat Insights* service is available. Red Hat Insights is a proactive service designed to enable you to identify, examine, and resolve known technical issues before they affect your deployment. Insights leverages the combined knowledge of Red Hat Support Engineers, documented solutions, and resolved issues to deliver relevant, actionable information to system administrators.

The service is hosted and delivered through the customer portal at <https://access.redhat.com/insights/> or through Red Hat Satellite. To register your systems, follow the [Getting Started Guide for Insights](#). For further information, data security and limits, refer to <https://access.redhat.com/insights/splash/>.

Red Hat Access Labs

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

- [Kickstart Configurator](#)

- » [Registration Assistant](#)
- » [NFS Helper](#)
- » [Linter for Dockerfile](#)
- » [Multipath Helper](#)
- » [iSCSI Helper](#)
- » [Code Browser](#)

Part I. New Features

This part describes new features and major enhancements introduced in Red Hat Enterprise Linux 6.9.

Chapter 2. General Updates

In-place upgrade from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7

An in-place upgrade offers a way of upgrading a system to a new major release of Red Hat Enterprise Linux by replacing the existing operating system. To perform an in-place upgrade, use the **Preupgrade Assistant**, a utility that checks the system for upgrade issues before running the actual upgrade, and that also provides additional scripts for the **Red Hat Upgrade Tool**. When you have solved all the problems reported by the **Preupgrade Assistant**, use the **Red Hat Upgrade Tool** to upgrade the system.

For details regarding procedures and supported scenarios, see the [Migration Planning Guide](#) and the [solution document dedicated to the upgrade](#).

Note that the **Preupgrade Assistant** and the **Red Hat Upgrade Tool** are available in the [Extras channel](#).

preupgrade-assistant rebased to version 2.3.3

The *preupgrade-assistant* packages have been upgraded to version 2.3.3, which provides a number of bug fixes, enhancements, and other changes over the previous version. Notably:

- ✦ A new **preupg -diff** tool has been added, which compares multiple Preupgrade Assistant XML reports: one new with unidentified problems and other reports with already analyzed problems. The tool helps to find issues that emerged in the new report by filtering out results that are the same in the new report and in at least one of the analyzed XML files. The output of the trimmed report is available in the XML and HTML format.
- ✦ Two new return codes have been added: **29** for **internal error**, and **30** for **user abort**.
- ✦ The meaning of the return code **22** has been changed to **invalid CLI option**.
- ✦ The STDOUT and STDERR output in the assessment report of the Preupgrade Assistant have been separated into two fields: **Additional output** for STDOUT and **Logs** for STDERR.
- ✦ The **python** module to be imported by the Preupgrade Assistant modules written in Python has been renamed from **preup** to **preupg**. Additionally, the **preup-ui-manage** executable has been renamed to **preupg-ui-manage**.
- ✦ The **exit_unknown** function and the **\$RESULT_UNKNOWN** variable have been removed. Instead of the **unknown** result, set the error result by using the **exit_error** function.
- ✦ The **set_component** module API function has been removed.
- ✦ The **component** input parameter has been removed from the following module API functions: **log_error**, **log_warning**, **log_info**, and **log_debug**. (BZ#[1427713](#), BZ#[1418697](#), BZ#[1392901](#), BZ#1393080, BZ#1372100, BZ#1372871)

Preupgrade Assistant enables blacklisting to improve performance

Preupgrade Assistant now supports creation of a blacklist file, which enables to skip all executable files on a path with a listed prefix. Users can activate this functionality in the `/etc/preupgrade-assistant.conf` file by setting the **exclude_file** value to the blacklist file name in the **xccdf_preupg_rule_system_BinariesRebuild_check** section. For example:

```
[xccdf_preupg_rule_system_BinariesRebuild_check]
exclude_file=/etc/pa_blacklist
```

Each line of the blacklist file contains a path prefix of executable files to be excluded. Previously, significant performance problems occurred when a large partition was mounted and the **RHEL6_7/system/BinariesRebuild** module checked numerous files on a list of executables. Now, users can filter out unimportant executable files and thus reduce time the module consumes. Note that this feature is expected to be changed in the future. (BZ#1392018)

Key file names unified in Preupgrade Assistant modules

Previously, each module in Preupgrade Assistant used different file names for certain required files, which made testing and orientation complicated. With this update, the key file names have been unified to **module.ini** (the metadata INI file), **check** (the check script), and **solution.txt** (a solution text) in each of the modules. Additionally, multiple rules (module IDs) have been renamed to conform with this change, so each rule now contains the unified **_check** suffix, for example, in the **result.html** and **result.xml** files. (BZ#1402478)

A new RHDS module to check a possibility of an in-place upgrade of an RHDS system

This update introduces a new Red Hat Directory Server (RHDS) module, which checks for relevant installed RHDS packages and gives users information about the possibility of an in-place upgrade of the RHDS system. As a result, if the relevant packages are installed, and the basic directory instance has been configured, the module creates a backup of the configuration files and prints information about them. (BZ#1406464)

cloud-init moved to the Base channel

As of Red Hat Enterprise Linux 6.9, the *cloud-init* package and its dependencies have been moved from the Red Hat Common channel to the Base channel. **Cloud-init** is a tool that handles early initialization of a system using metadata provided by the environment. It is typically used to configure servers booting in a cloud environment, such as OpenStack or Amazon Web Services. Note that the *cloud-init* package has not been updated since the latest version provided through the Red Hat Common channel. (BZ#1421281)

Chapter 3. Authentication and Interoperability

SSSD now enables the administrator to select which domains from the AD forest can be contacted

In some environments, only a subset of domains in a joined Active Directory (AD) forest can be reached. Attempting to contact an unreachable domain might cause unwanted timeouts or switch the System Security Services Daemon (SSSD) to offline mode.

To prevent this, the administrator can now configure a list of domains to which SSSD connects by setting the `ad_enabled_domains` option in the `/etc/sss/sss.conf/` file. For details, see the `sss-ad(5)` man page. (BZ#1324428)

SSSD now enables selecting a list of PAM services that will not receive any environmental variables from `pam_sss`

In some cases, it is not desirable to propagate environment variables set by the `pam_sss` Pluggable Authentication Module (PAM). For example, when using the `sudo -i` command, users might want to transfer the `KRB5CCNAME` variable of the original user to the target environment.

Previously, when a non-privileged user executed the `sudo -i` command to become another non-privileged user, the new non-privileged user did not have the permissions to read the Kerberos credentials cache that `KRB5CCNAME` pointed to.

For this use case, this update adds a new option named `pam_response_filter`. Using `pam_response_filter`, the administrator can list PAM services (such as `sudo -i`) that do not receive any environmental variables (such as `KRB5CCNAME`) during login. Now, if `pam_response_filter` lists `sudo -i`, a user can switch from one non-privileged user to another without `KRB5CCNAME` being set in the target environment. (BZ#1329378)

IdM servers can now be configured to require TLS 1.2 or better

Version 1.2 of the Transport Layer Security (TLS) protocol is considered significantly more secure than previous versions. This update enables you to configure your Identity Management (IdM) server to forbid communication using protocols that are less secure than `TLS 1.2`.

For details, see the following Red Hat Knowledgebase article:

<https://access.redhat.com/articles/2801181>. (BZ#1367026)

`pam_faillock` can be now configured with `unlock_time=never`

The `pam_faillock` module now allows specifying using the `unlock_time=never` option that the user authentication lock caused by multiple authentication failures should never expire. (BZ#1404832)

The `libkadm5*` libraries have been moved to the `libkadm5` package

In Red Hat Enterprise Linux 6.9, the `libkadm5*` libraries have been moved from the `krb5-libs` to the new `libkadm5` package. As a consequence, `yum` is not able to downgrade the `krb5-libs` package automatically. Before downgrading, remove the `libkadm5` package manually:

```
# rpm -e --nodeps libkadm5
```

After you have manually removed the package, use the **yum downgrade** command to downgrade the *krb5-libs* package to a previous version. (BZ#1351284)

Chapter 4. Clustering

Support added for Oracle 11g in Oracle and OrLsnr Pacemaker resource agents

As of Red Hat Enterprise Linux release 6.9, the Pacemaker resource agents **Oracle** and **OrLsnr** support Oracle database 11g. (BZ#1336846)

Pacemaker now supports alert agents

You can now create **Pacemaker** alert agents to take some external action when a cluster event occurs. The cluster passes information about the event to the agent by means of environment variables. Agents can do anything desired with this information, such as send an email message, log to a file, or update a monitoring system. For information on configuring alert agents, see Configuring the Red Hat High Availability Add-On with Pacemaker. (BZ#[1253325](#), BZ#[1376480](#))

clufter is now fully supported

The *clufter* packages provide a tool for transforming and analyzing cluster configuration formats. They can be used to assist with migration from an older stack configuration to a newer configuration that leverages Pacemaker. The **clufter** tool, previously available as a Technology Preview, is now fully supported. For information on the capabilities of **clufter**, see the **clufter(1)** man page or the output of the **clufter -h** command. For examples of **clufter** usage, see the following Red Hat Knowledgebase article: <https://access.redhat.com/articles/2810031>. (BZ#1318326)

clufter rebased to version 0.59.8

The *clufter* packages have been upgraded to upstream version 0.59.8, which provides a number of bug fixes, new features, and user experience enhancements over the previous version. Among the notable updates are the following:

- ✦ When converting either CMAN or Pacemaker stack specific configuration into the respective sequence of **pcs** commands with the ***2pcscmd** families of commands, the **clufter** tool no longer suggests **pcs cluster cib file --config**, which does not currently work for subsequent local-modification **pcs** commands. Instead it suggests **pcs cluster cib file**. (RHBZ#1328078)
- ✦ The **clufter** tool outputs now may vary significantly depending on the specified distribution target since the tool now aligns the output with what the respective environment, such as the **pcs** version, can support. Because of this, your distribution or setup may not be supported, and you should not expect that one sequence of **pcs** commands that the **clufter** tool produces is portable to a completely different environment.
- ✦ The **clufter** tool now supports several new features of the **pcs** tool, including alert handlers configuration. Additionally, the **clufter** tool supports older features recently added to the **pcs** tool, including resource sets for colocation and order constraints.
- ✦ When converting either CMAN + RGManager stack specific configuration into the respective Pacemaker configuration (or sequence of **pcs** commands reflecting the same) with the **ccs2pcs*** families of commands, the **clufter** tool no longer refuses to convert entirely valid lvm resource agent configuration, which could happen before. (BZ#[1367536](#))

luci interface allows administrators to verify authenticity of remote machines

An encrypted channel requires established authenticity between the endpoints to be reasonably secure and protected against man-in-the-middle attacks. Administrators using **luci** to manage clusters are now automatically provided with the corresponding certificate fingerprints of cluster nodes that are entered when creating a new cluster, adding nodes to a cluster, or adding an existing cluster to ***luci***'s management. This allows administrators to verify the authenticity of remote machines first before entrusting the remote nodes with credentials during standard, inverse (self-against-remote) authentication. (BZ#885028)

luci now lists explicit configured actions for individual resources

In a cluster configuration, it is useful to be able to review configured actions for given resources. This may be particularly useful when verifying that the implicit operations, such as the `'depth'` parameter of the **status** action, are overwritten with user configuration. More generally, being able to review configured actions can show the affect that these modifications and additions to implicit actions have on current cluster behavior.

luci now lists configured actions per individual resources in the **Service Groups** breakdown view, showing which parameters are disregarded for particular actions and emphasizing timeouts if they are set as enforced. Note that the view does not allow for active modifications of the actions; to modify the actions, you use the `--addaction` and `--rmaction` parameters of the **ccs** CLI tool. (BZ#[1173942](#))

Chapter 5. Compiler and Tools

Support for the `el_GR@euro`, `ur_IN`, and `wal_ET` locales has been added

The `el_GR@euro`, `ur_IN`, and `wal_ET` locales provide specialized support for newer currency symbols like the Euro, and complete coverage in the instances where the locale was previously unsupported.

Users can now specify these locales using the relevant environment variables to take advantage of the new localization support. (BZ#[1101858](#))

The `Net::SSLeay` Perl module now supports restricting of TLS version

The `Net::SSLeay` Perl module has been updated to support explicit specification of the TLS protocol version, which can be used for improving security. To restrict TLS version to 1.1 or 1.2, set the `Net::SSLeay::ssl_version` variable to `11` or `12`, respectively. (BZ#[1325407](#))

The `IO::Socket::SSL` Perl module now supports restricting of TLS version

The `Net::SSLeay` Perl module has been updated to support explicit specification of the TLS protocol versions 1.1 or 1.2 to improve security, and the `IO::Socket::SSL` module has been updated accordingly. When a new `IO::Socket::SSL` object is created, it is now possible to restrict the TLS version to 1.1 or 1.2 by setting the `SSL_version` option to `TLSv1_1` or `TLSv1_2` respectively. Alternatively, `TLSv11` and `TLSv12` can be used. Note that these values are case-sensitive. (BZ#[1331037](#))

`ca-certificates` rebased to version 2.10

The certificate store has been upgraded to include the changes contained in version 2.10 of the Certificate Authority certificate list published by the Mozilla Foundation as part of the Network Security Services (NSS) version 3.27. In order to preserve compatibility with existing PKI deployments and with software based on OpenSSL and GnuTLS, several root CA certificates with an RSA key size of 1024 bits have been kept as trusted by default. See the following Knowledgebase article for instructions on disabling these legacy modifications: <https://access.redhat.com/articles/1413643>. (BZ#[1368996](#))

Chapter 6. Directory Server in Red Hat Enterprise Linux

Directory Server now supports enabling and disabling specific TLS versions

Previously, Directory Server running on Red Hat Enterprise Linux 6 provided no configuration options to enable or disable specific TLS versions. For example, it was not possible to disable the insecure TLS 1.0 protocol while keeping later versions enabled. This update adds the **nsTLS10**, **nsTLS11**, and **nsTLS12** parameters to the **cn=encryption, cn=config** entry. As a result, it is now possible to configure specific TLS protocol versions in Directory Server.

Note, that these parameters have a higher priority than the **nsTLS1** parameter, that enables or disables all TLS protocol versions. (BZ#[1330758](#))

Chapter 7. Hardware Enablement

cpuid is now available

With this update, the **cpuid** utility is available in Red Hat Enterprise Linux. This utility dumps detailed information about the CPU(s) gathered from the CPUID instruction, and also determines the exact model of CPU(s). It supports Intel, AMD, and VIA CPUs. (BZ#1316998)

Support for RealTek RTS5250S SD4.0 Controllers

The Realtek RTS5205 card reader controllers have been added to the kernel. (BZ#1167938)

Chapter 8. Installation and Booting

The `NO_DHCP_HOSTNAME` option has been added

The `NO_DHCP_HOSTNAME` option can now be specified in the `/etc/sysconfig/network` configuration file. Previously, in certain situations it was not possible to prevent initialization scripts from obtaining the host name through DHCP, even when using a static configuration. With this update, if the `NO_DHCP_HOSTNAME` option is set to **yes**, **true**, or **1** in the `/etc/sysconfig/network` file, initialization scripts are prevented from obtaining the host name through DHCP. (BZ#[1157856](#))

Chapter 9. Kernel

Chelsio firmware updated to version 1.15.37.0

Chelsio firmware has been updated to version 1.15.37.0, which provides a number of bug fixes and enhancements over the previous version.

The most notable bug fixes are:

- The **iscsi tlv** driver is no longer incorrectly sent to host.
- The firmware no longer terminates unexpectedly due to enabling or disabling the Data Center Bridging Capability Exchange (DCBX) protocol.
- The app priority value is now handled correctly in the firmware. (BZ#1349112)

The **bnxt_en** driver updated to the latest upstream version

The **bnxt_en** driver has been updated with several minor fixes and with support for BCM5731X, BCM5741X, and 57404 Network Partitioning (NPAR) devices. (BZ#1347825)

The **ahci** driver supports Marvell 88SE9230

The **ahci** driver now supports Marvell 88SE9230 controller. (BZ#1392941)

Chapter 10. Networking

NetworkManager now supports manual DNS configuration with `dns=none`

With this update, the user has the option to prevent **NetworkManager** from modifying the `/etc/resolv.conf` file. This is useful for manual management of DNS settings. To protect the file from being modified, add the `dns=none` option to the `/etc/NetworkManager/NetworkManager.conf` file. (BZ#1308730)

Chapter 11. Security

TLS 1.2 support added to all system components

With the addition of **TLS 1.2** support to the **GnuTLS** component, Red Hat Enterprise Linux 6 offers complete support for **TLS 1.2** in the shipped security libraries: **OpenSSL**, **NSS**, and **GnuTLS**. Several modern standards such as PCI-DSS v3.1 recommend the latest **TLS** protocol, which is currently **TLS 1.2**. This addition allows you to use Red Hat Enterprise Linux 6 with future revisions of security standards, which may require **TLS 1.2** support.

For more information about the cryptographic changes in the Red Hat Enterprise Linux 6, see this article on the Red Hat Customer Portal: <https://access.redhat.com/blogs/766093/posts/2787271>. (BZ#1339222)

OpenSCAP 1.2.13 is NIST certified

OpenSCAP 1.2.13 has been certified by the National Institute of Standards and Technology's (NIST) Security Content Automation Protocol (SCAP) 1.2 in the Authenticated Configuration Scanner category with the Common Vulnerabilities and Exposure (CVE) option. **OpenSCAP** provides a library that can parse and evaluate each component of the SCAP standard. This makes creating new SCAP tools convenient. Also, **OpenSCAP** offers a multi-purpose tool designed to format content into documents or scan a system based on this content. (BZ#1364207)

vsftpd now uses TLS 1.2 by default

Users of the Very Secure File Transfer Protocol (FTP) daemon (*vsftpd*) can select a specific version of **TLS** protocol up to 1.2. **TLS 1.2** has been enabled by default to bring security of *vsftpd* to the same level as the same package in Red Hat Enterprise Linux 7. New default ciphers specific to **TLS 1.2** has been added: **ECDHE-RSA-AES256-GCM-SHA384** and **ECDHE-ECDSA-AES256-GCM-SHA384**. These changes do not break existing configurations. (BZ#1350724)

auditd now supports incremental_async

The **auditd** daemon now supports a new flush technique called **incremental_async**. This new mode significantly improves the **auditd** daemon's logging performance maintaining short flush intervals for security. (BZ#1369249)

scap-security-guide now supports ComputeNode

The *scap-security-guide* project now supports scanning of the ComputeNode variant of Red Hat Enterprise Linux and the *scap-security-guide* package is also distributed in the relevant channel. (BZ#[1311491](#))

rsyslog7 now enables TLS 1.2

With this update, the **rsyslog7** multi-threaded syslog daemon explicitly enables **TLS 1.2** in the **GnuTLS** component. (BZ#1323199)

Chapter 12. Servers and Services

A DHCP client hook example added for DDNS for Microsoft Azure cloud

An example of the **DHCP** client hook for Dynamic DNS (DDNS) for Microsoft Azure cloud has been added to the *dhcp* package. The administrator can now easily enable this hook, and register Red Hat Enterprise Linux clients with a **DDNS** server. (BZ#1321945)

postfix now supports user-controlled configuration of TLS

With this update, *postfix* offers configuration options for more precise control of the Transport Layer Security (TLS) protocol version. For example, you can now disable **TLS** v1.1 while having **TLS** v1.2 enabled. To do this, add the following line to the **main.cf** file:

```
smtpd_tls_mandatory_protocols = !TLSv1.1
```

(BZ#[1287192](#))

Chapter 13. Storage

The smartPQI (`smartpqi`) driver is now available

This update provides the smartPQI (`smartpqi`) driver for new Microsemi storage adapter hardware, which becomes available in 2017. The new hardware can also be used with the previous `aacraid` driver on Red Hat Enterprise Linux 6.5, 6.6, 6.7, and 6.8. In comparison with the `aacraid` driver, the `smartpqi` driver provides improved performance and enhanced functionality.

Migration from Red Hat Enterprise Linux 6.8 to Red Hat Enterprise Linux 6.9 changes the driver from `aacraid` to `smartpqi`. As long as standard installation configurations are used, this driver change is transparent to the user and no action is needed. The new `smartpqi` driver is automatically used after booting Red Hat Enterprise Linux 6.9. (BZ#1343743)

Update of `mpt3sas`

The `mpt3sas` storage driver has been updated to version 14.100.00.00-rh, which adds support for new devices with these PCI IDs:

- 0x1000:0x00AA
- 0x1000:0x00AB SAS3516 Fusion-MPT Tri-Mode RAID On Chip (ROC)
- 0x1000:0x00AC SAS3416 Fusion-MPT Tri-Mode I/O Controller Chip (IOC)
- 0x1000:0x00AD
- 0x1000:0x00AE SAS3508 Fusion-MPT Tri-Mode RAID On Chip (ROC)
- 0x1000:0x00AF SAS3408 Fusion-MPT Tri-Mode I/O Controller Chip (IOC) (BZ#1306469)

Update of `megaraid_sas`

The `megaraid_sas` driver has been updated to version 07.700.00.00-rc1, which adds support for new devices with these PCI IDs:

- 0x1000:0x0014
- 0x1000:0x0016
- 0x1000:0x0017
- 0x1000:0x001B
- 0x1000:0x001C (BZ#1306457)

A new default configuration for Huawei XSG1 arrays has been added for `device-mapper-multipath`

On Red Hat Enterprise Linux 6, a specific configuration is recommended in the `device-mapper-multipath` tool configuration for Huawei XSG1 arrays. This configuration is now used by default. (BZ#1333334)

The `disable_changed_wwids multipath.conf` option is now available in `multipath` to avoid data corruption

The multipath tool now has the `disable_changed_wwids` `multipath.conf` option. If `disable_changed_wwids` is set to `yes`, the `multipathd` service monitors path devices, and if their World Wide Identifier (WWID) changes, `multipathd` disables access to the path devices until the WWID changes back.

If a Logical Unit Number (LUN) is remapped while a multipath device exists on top of it, it is possible in some cases for I/O to be written to an incorrect LUN, which leads to corruption. Writing to an incorrect LUN can be detected by `multipathd` that registers a change of the LUN WWID, and disables access to the device.

Note that due to the gap between when the LUN is remapped, and when `multipathd` is notified that the device has changed, there is still a risk of corruption in some cases, and remapping in-use LUNs is still not supported. (BZ#[1377532](#))

device-mapper-multipath now supports the `max_sectors_kb` configuration parameter

With this update, `device-mapper-multipath` provides a new `max_sectors_kb` parameter in the defaults, devices, and multipaths sections of the `multipath.conf` file. The `max_sectors_kb` parameter allows you to set the `'max_sectors_kb` device queue parameter to the specified value on all underlying paths of a multipath device before the multipath device is first activated.

When a multipath device is created, the device inherits the `max_sectors_kb` value from the path devices. Manually raising this value for the multipath device or lowering this value for the path devices can cause multipath to create I/O operations larger than the path devices allow.

Using the `max_sectors_kb` `multipath.conf` parameter is an easy way to set these values before a multipath device is created on top of the path devices, and prevent invalid-sized I/O operations from being passed down. (BZ#[1355669](#))

The `skip_kpartx` `multipath.conf` option to allow skipping `kpartx` partition creation has been added

This update enables the user to only create a multipath device, and not any partitions, even if the device has a partition table. Now, multipath devices that are configured with the `skip_kpartx` option do not have any partition devices created for them. (BZ#[1310320](#))

Users are now warned if they create multipath devices while `multipathd` is not running

With this update, multipath prints a warning message for adding and listing multipath devices when the `multipathd` service is not running. (BZ#[1305589](#))

Chapter 14. Virtualization

Configuration options can be used to exclude weak ciphers

Previously, *libvirt* depended on the hard-coded cipher defaults in **GnuTLS**. This made it possible to use weak ciphers. With this update, configuration options to exclude weak ciphers have been added to the **libvirtd.conf** and **libvirt.conf** files. In addition, **TLS** priority support was added to *libvirt* URIs. As a result, the list of used ciphers can be customized to exclude weak ciphers.

(BZ#[1333415](#))

Improved Hyper-V storage driver performance

The *storvsc* Hyper-V storage driver was updated from upstream. This provides moderate performance improvement of I/O operations when using the Hyper-V *storvsc* driver for certain workloads.

(BZ#1352824)

Hyper-V clock source changed to use the TSC page

With this update, the Time Stamp Counter (TSC) page is used as the Hyper-V clock source. The TSC page provides a more efficient way of computing the per-guest reference counter value than the previously used model-specific register (MSR). As a result, kernel operations that involve reading time stamps are now faster.

Note that this feature is only supported on 64-bit kernels. (BZ#1365049)

Setting the account password is now possible for any guest user

The **guest-set-user-password** command has been introduced for the QEMU guest agent. This allows setting the account password for any guest user, including the root, when using QEMU and KVM. (BZ#[1303906](#))

Chapter 15. Red Hat Software Collections

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

Red Hat Developer Toolset is designed for developers working on the Red Hat Enterprise Linux platform. It provides current versions of the GNU Compiler Collection, GNU Debugger, and other development, debugging, and performance monitoring tools. Since Red Hat Software Collections 2.3, the Eclipse development platform is provided as a separate Software Collection.

Dynamic languages, database servers, and other tools distributed with Red Hat Software Collections do not replace the default system tools provided with Red Hat Enterprise Linux, nor are they used in preference to these tools. Red Hat Software Collections uses an alternative packaging mechanism based on the `sc1` utility to provide a parallel set of packages. This set enables optional use of alternative package versions on Red Hat Enterprise Linux. By using the `sc1` utility, users can choose which package version they want to run at any time.



Important

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

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

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

Part II. Known Issues

This part documents known problems in Red Hat Enterprise Linux 6.9.

Chapter 16. General Updates

The default value of `first_valid_uid` in Dovecot has changed in Red Hat Enterprise Linux 7

Since Red Hat Enterprise Linux 7.3, the default value of the `first_valid_uid` configuration option of **Dovecot** has changed from **500** in Red Hat Enterprise Linux 6 to **1000** in Red Hat Enterprise Linux 7. Consequently, if a Red Hat Enterprise Linux 6 installation does not have `first_valid_uid` explicitly defined, the **Dovecot** configuration will not allow users with UID less than **1000** to log in after the update to Red Hat Enterprise Linux 7.

To avoid breaking the configuration, redefine `first_valid_uid` to **500** after the upgrade in the `/etc/dovecot/conf.d/10-mail.conf` file. Note that only installations where `first_valid_uid` is not explicitly defined are affected by this problem. (BZ#[1388967](#))

Incorrect information about the expected default settings of services in Red Hat Enterprise Linux 7

The module of Preupgrade Assistant that handles initscripts provides incorrect information about the expected default settings of the services in Red Hat Enterprise Linux 7 according to the `/usr/lib/systemd/system-preset/90-default.preset` file in Red Hat Enterprise Linux 7 and according to the current settings of the Red Hat Enterprise Linux 6 system. In addition, the module does not check the default settings of the system but only the settings for the runlevel used during the processing of the check script, which might not be the default runlevel of the system. As a consequence, initscripts are not handled in the anticipated way and the new system needs more manual action than expected. However, the user is informed about the settings that will be chosen for relevant services, despite the presumable default settings. (BZ#1366671)

Chapter 17. Authentication and Interoperability

SSSD fails to manage sudo rules from the IdM LDAP tree

The System Security Services Daemon (SSSD) currently uses the IdM LDAP tree by default. As a consequence, it is not possible to assign sudo rules to non-POSIX groups. To work around this problem, modify the `/etc/sss/sss.conf` file to set your domain to use the **compat** tree again:

```
[domain/EXAMPLE]
...
ldap_sudo_search_base = ou=sudoers,dc=example,dc=com
```

As a result, SSSD will load sudo rules from the **compat** tree and you will be able to assign rules to non-POSIX groups.

Note that Red Hat recommends to configure groups referenced in sudo rules as POSIX groups. (BZ#1336548)

winbindd crashes when installing a new AD trust

When configuring a new Active Directory (AD) trust on a newly installed system, the `ipa-adtrust-install` utility might report that the `winbindd` service terminated unexpectedly. Otherwise, `ipa-adtrust-install` completes successfully.

If this problem occurs, restart the IdM services by using the `ipactl restart` command after running `ipa-adtrust-install`. This also restarts `winbindd`.

Note that the full extent of the functional impact of this problem is still unknown. Some trust functionality might not work until `winbindd` is restarted. (BZ#[1399058](#))

ns1cd fails to resolve user or group identities when it is started before the network connection is fully up

When `ns1cd`, the local LDAP name service daemon, is started before the network connection is fully up, the daemon fails to connect to an LDAP server. As a consequence, resolving user or group identities does not work. To work around this problem, start `ns1cd` after the network connection is up. (BZ#[1401632](#))

Chapter 18. Desktop

The vmware driver does not support multiple displays

The **vmware** video driver for the X11 window system misses certain features related to multi-display support. As a consequence, Red Hat Enterprise Linux 6 guests running on VMware cannot correctly use multiple displays and only single-display support is available.

Please contact Red Hat support for test packages if you require multi-display support. (BZ#[1320480](#))

Incorrect mouse pointer movement after screen rotation inside a virtual machine in VMWare 11 or VMWare 12

If the screen rotation is changed inside a virtual machine in VMWare 11 or VMWare 12, the pointer movement remains unchanged. This only happens when the **xorg-x11-drv-vmware** driver is used, which initializes an absolute-axis device rather than a relative-axis device. The pointer does not follow the expected path because the driver is still mapping to the original coordinate system. To work around this problem, it is necessary to manually rotate the device, for example by running the following command:

```
xinput set-prop "ImPS/2 Generic Wheel Mouse" "Coordinate Transformation Matrix" 0 -1 1 1 0 0 0 0 1
```

Note that the command above is only an example. In general, the matrix needs to be adjusted depending on the specific scenario. Once the matrix is applied, pointer movement matches the rotation of the screen. (BZ#[1322712](#), BZ#[1318340](#))

Using Radeon or Nouveau can cause incorrectly rendered graphics

A bug in the Xorg server can, under rare circumstances, cause graphics to be rendered incorrectly if using the Radeon or Nouveau graphics device driver. For example, the Thunderbird message pane can be displayed incorrectly.

For Nouveau, as a workaround, add the **WrappedFB** option to the **xorg.conf** file as follows:

```
Section "Device"
    Identifier "nouveau-device"
    Driver "nouveau"
    Option "WrappedFB" "true"
EndSection
```

This workaround avoids the faulty logic in the X server, and the Thunderbird message pane will be displayed correctly. (BZ#[1076595](#))

Chapter 19. Directory Server in Red Hat Enterprise Linux

IdM schema replications from Red Hat Enterprise Linux 7 to 6.9 fail

Identity Management (IdM) in Red Hat Enterprise Linux 6.9 uses a different schema definition in the **nsEncryptionConfig** object class than IdM on Red Hat Enterprise Linux 7.3. Because the schema learning mechanism is unable to merge definitions, schema replications between servers fail. As a consequence, mechanisms relying on the schema can fail. For example, schema violations and plug-in failures can occur, replication can fail, and access control instructions (ACI) can be ignored. In an upcoming Red Hat Enterprise Linux 7.3 update, the **nsTLS10**, **nsTLS11**, and **nsTLS12** attributes will be added to the list of allowed attributes in the **nsEncryptionConfig** object class, and as a consequence, mechanisms relying on the schema no longer fails in the described scenario. (BZ#[1404443](#))

Chapter 20. Installation and Booting

The installer displays the number of multipath devices, and number of multipath devices selected, incorrectly

Multipath devices are configured properly, but the installer displays the number of devices and number of selected devices incorrectly. There is no known workaround at this point. (BZ#914637)

The installer displays the amount of disk space within multipath devices incorrectly

Multipath devices are configured properly, but the installer displays disk space and number of devices incorrectly. There is no known workaround at this point. (BZ#1014425)

The `device.map` configuration file generated by Anaconda is sometimes incorrect

Due to limitations in the kernel, the `device.map` configuration file that is used to map BIOS drives to operating system devices might be generated incorrectly in certain situations, particularly when installing from a USB key. As a consequence, booting sometimes fails after installation. To work around this problem, manually update the `device.map` file in the `/boot/grub` directory. After updating `device.map` so that it correctly maps devices on the system, Red Hat Enterprise Linux 6 will boot as expected. (BZ#[1253223](#))

The `ifup` script incorrectly replaces manually-defined default routes

If a default route is manually added to the routing table, the `ifup` script incorrectly replaces it, when setting up other interfaces, if the `GATEWAY` parameter is specified. To work around this bug, specify a non-zero metric for either the manually-added route, or when adding a route with `ifup`. (BZ#1090559)

Upgrading Red Hat Enterprise Linux 6 on UEFI systems clears the boot loader password

When upgrading Red Hat Enterprise Linux 6 on a system with UEFI firmware and a boot loader password set, the boot loader password is removed. As a consequence, modifying the boot record is possible without a password. To work around this problem, make a back up of the password settings from the `/boot/efi/EFI/redhat/grub.conf` configuration file before upgrading, and then restore the settings to the `/boot/efi/EFI/redhat/grub.conf` file in the new system. (BZ#1416653)

Chapter 21. Kernel

Certain NIC firmware can become unresponsive with the `bnx2x` driver

Due to a bug in the unload sequence of the pre-boot drivers, the firmware of some internet adapters can become unresponsive after the `bnx2x` driver takes over the device. The `bnx2x` driver detects the problem and returns the message in the kernel log:

```
Storm stats were not updated for 3 times.
```

To work around this problem, apply the latest NIC firmware updates provided by your hardware vendor. As a result, unloading of the pre-boot firmware now works as expected and the firmware no longer hangs after `bnx2x` takes over the device. (BZ#1012684)

e1000e cards might not get an IPv4 address

Some e1000e network interface cards (NICs) might fail to get an IPv4 address assigned after the system is rebooted. To work around this problem, add the following line to the `/etc/sysconfig/network-scripts/ifcfg-<interface>` file:

```
LINKDELAY=10
```

(BZ#822725)

The `ecb` kernel module fails when `dracut` is not upgraded

When upgrading only the kernel rpm from Red Hat Enterprise Linux 6.7 to version 6.8, upgrade the `dracut` package to the latest version (`dracut-004-409.el6.rpm`).

Upgrading `dracut` enables the `ecb` module to work. The `ecb` kernel module is needed by the `drbg` kernel module when using the Advanced Encryption Standard (AES) implementation on non-x86 architectures. If you do not upgrade `dracut`, the `drbg` AES implementation fails with a warning message, although other `drbg` modules still work. (BZ#1315832)

Chapter 22. Networking

The radvd occasionally terminates unexpectedly due to a race condition

In the **Router Advertisement Daemon** (radvd), there is a race condition in **radvd** timer handling. Consequently, the **radvd** occasionally terminates unexpectedly. (BZ#[1058698](#))

Chapter 23. Servers and Services

Printing a PDF file upside down with cups is currently impossible

In the CUPS printing system, the `-o orientation-requested=6` option in the `lp -d [printer] -o orientation-requested=6 [filename]` command, which is expected to rotate the printed page by 180°, does not work. (BZ#1099617)

Printing PDF files using the fit-to-page and fitplot options does not work on printers with hardware margins

In the CUPS printing system, the `lp -d printer-with-hwmargins -o fit-to-page` and `lp -d printer-with-hwmargins -o fitplot` commands use the `-o fit-to-page` and `-o fitplot` options which resize the document to be printed so that it fits the paper size. The options do not work for printing PDF files on printers with hardware margins. (BZ#1268131)

DHCP client sends unicast requests through the incorrect interface

DHCP client does not support multiple interfaces on the same subnet and it is not able to ensure that unicast requests go through the right interface. Consequently, **DHCP** client fails to renew a lease, and network configuration stops working. There is no known workaround at this point. **DHCP** client cannot be used in configuration with two interfaces connected to the same subnet. (BZ#[1297445](#))

A *.dsc file converted from a *.pdf file by the pdf2dsc script cannot be opened in Evince

It is no longer possible to convert a *.pdf (Portable Document Format) file into a *.dsc (Document Structure Convention) file with the `pdf2dsc` script, and open the converted *.dsc file with the Evince GNOME document viewer, located outside the Ghostscript's sandbox. It is a result of the fixed `-dSAFER` option, which forces Ghostscript to operate in sandbox mode. For details and a workaround, see <https://access.redhat.com/articles/2948831>. (BZ#[1411843](#))

Chapter 24. Storage

mkfs might become temporarily unresponsive on the 32-bit AMD and Intel architecture

Starting with Red Hat Enterprise Linux 6.9, the maximum amount of time the bdi forker thread sleeps for is increased. As a consequence, the **mkfs** command might become unresponsive for up to five minutes on the 32-bit AMD and Intel architecture. (BZ#1111683)

Chapter 25. System and Subscription Management

ReaR works only on the eth0 interface

ReaR produces a rescue system that does not support mounting an NFS server using an interface other than eth0. Consequently, the backup files cannot be downloaded and the system cannot be restored. To work around this problem, ensure that the used interface is eth0 by restarting dhclient. (BZ#[1313417](#))

ReaR creates two ISO images instead of one

In ReaR, the **OUTPUT_URL** directive enables specifying location for the ISO image containing the rescue system. Currently, with this directive set, ReaR creates two copies of the ISO image: one in the specified directory and one in the **/var/lib/rear/output/** default directory. This requires additional space for the image. This is especially important if a full-system backup is included into the ISO image (using the **BACKUP=NETFS** and **BACKUP_URL=iso:///backup/** configuration).

To work around this behavior, delete the extra ISO image once ReaR has finished working or, to avoid having a period of time with double storage consumption, create the image in the default directory and then move it to the desired location manually.

There is a request for enhancement to change this behavior and make ReaR create only one copy of the ISO image. (BZ#1320551)

Chapter 26. Virtualization

coolkey does not load on Windows 7 guests

Loading the **Coolkey** module on Windows 7 guest virtual machines currently fails, which prevents smart card redirection from working properly on these guests. (BZ#[1331471](#))

Disabling vCPUs on Hyper-V guests fails

Currently, it is not possible to disable CPUs on guest virtual machines running on Microsoft Hyper-V, including Microsoft Azure cloud, due to the lack of support from the host side. However, it is possible to reduce the number of online CPUs by booting guests with the **nr_cpus=XX** parameter passed on the kernel command line, where **XX** is the number of online CPUs required.

For more information, see <https://access.redhat.com/solutions/2790331>. (BZ#1396336)

Hot plugging hard disks as a batch on the VMware ESXi hypervisor does not work reliably

When hot plugging multiple hard disks at the same time to a Red Hat Enterprise Linux 6 guest virtual machine running on the VMware ESXi hypervisor, the host currently does not inform the guest about all of the added disks, and some of the disks thus cannot be used. To work around this problem, hot plug one hard disk at a time in the described scenario. (BZ#1224673)

Guests cannot access floppy disks larger than 1.44 MB

Guest virtual machines are currently unable to access floppy drive images larger than 1.44 MB if they are inserted while the guest is running. To work around the problem, insert the floppy drive image prior to booting the guest. (BZ#[1209362](#))

Hyper-V guest integration services stop working after they are disabled and re-enabled

Currently, Red Hat Enterprise Linux 6 guest virtual machines running on the Microsoft Hyper-V hypervisor do not automatically restart the **hyperv-daemons** suite after Hyper-V guest integration services, such as data exchange and backup, are disabled and then re-enabled. As a consequence, these integration interface services stop working after they are disabled and re-enabled in the Hyper-V Manager interface.

To work around this problem, restart the **hypervkvpd**, **hypervvssd**, and **hypervfcopyd** services in the guest after re-enabling the integration services from Hyper-V Manager, or do not change the status of the integration services while the guest is running. (BZ#1121888)

Booting virtual machines with the fsgsbase and smep flags on older host CPUs fails

The **fsgsbase** and **smeop** CPU flags are not properly emulated on certain older CPU models, such as the early Intel Xeon E processors. As a consequence, using **fsgsbase** or **smeop** when booting a guest virtual machine on a host with such a CPU causes the boot to fail. To work around this problem, do not use **fsgsbase** and **smeop** if the CPU does not support them. (BZ#[1371765](#))

Guests with recent Windows systems in some cases fail to boot if `hv_relaxed` is used

Attempting to boot KVM guests with the following operating systems currently fails with an **error code: 0x0000001E** message if the value of the `-cpu` option is **SandyBridge** or **Opteron_G4** and the `hv_relaxed` option is used.

- 64-bit Windows 8 or later
- 64-bit Windows Server 2012 or later

To work around this problem, do not use `hv_relaxed`. (BZ#[1063124](#))

Limited CPU support for Windows 10 and Windows Server 2016 guests

On a Red Hat Enterprise 6 host, Windows 10 and Windows Server 2016 guests can only be created when using the following CPU models:

- the Intel Xeon E series
- the Intel Xeon E7 family
- Intel Xeon v2, v3, and v4
- Opteron G2, G3, G4, G5, and G6

For these CPU models, also make sure to set the CPU model of the guest to match the CPU model detected by running the `virsh capabilities` command on the host. Using the application default or hypervisor default prevents the guests from booting properly.

To be able to use Windows 10 guests on Legacy Intel Core 2 processors (also known as Penryn) or Intel Xeon 55xx and 75xx processor families (also known as Nehalem), add the following flag to the Domain XML file, with either Penryn or Nehalem as MODELNAME:

```
<cpu mode='custom' match='exact'>
  <model>MODELNAME</model>
  <feature name='erms' policy='require' />
</cpu>
```

Other CPU models are not supported, and both Windows 10 guests and Windows Server 2016 guests created on them are likely to become unresponsive during the boot process. (BZ#[1346153](#))

Windows 8 Virtual Machine Now Shuts Down Properly

Previously, when powering down a Windows 8 guest virtual machine with an i386 processor, the virtual machine got stuck on the Shutting Down screen. With the update of `virtio` and `qxl` drivers, this has been fixed. Consequently, the shutdown process finishes. (BZ#[1271469](#))

Network connectivity not restored when vnic is enabled

If the `netdev(tap)` link is set to off and the `vnic(virtio-net/e1000)` link is set to on, network connectivity does not resume. However, if the `vnic(virtio-net/e1000)` link is set to off and the `netdev(tap)` link is set to on, network connectivity resumes.

To resolve the issue, consistently use the same device to control the link. If `netdev(tap)` link was set to off, using it to turn the link back on will work correctly. (BZ#1198956)

Appendix A. Component Versions

This appendix is a list of components and their versions in the Red Hat Enterprise Linux 6.9 release.

Table A.1. Component Versions

Component	Version
Kernel	2.6.32-696
QLogic q1a2xxx driver	8.07.00.26.06.8-k
QLogic ql2xxx firmware	<i>ql2100-firmware-1.19.38-3.1</i>
	<i>ql2200-firmware-2.02.08-3.1</i>
	<i>ql23xx-firmware-3.03.27-3.1</i>
	<i>ql2400-firmware-7.03.00-1</i>
	<i>ql2500-firmware-7.03.00-1</i>
Emulex lpfc driver	0:11.0.0.5
iSCSI initiator utils	<i>iscsi-initiator-utils-6.2.0.873-26</i>
DM-Multipath	<i>device-mapper-multipath-0.4.9-100</i>
LVM	<i>lvm2-2.02.143-12</i>

Appendix B. Revision History

Revision 0.0-9	Tue Mar 28 2017	Lenka Špačková
Added a feature and three known issues to Virtualization.		
Revision 0.0-8	Tue Mar 21 2017	Lenka Špačková
Release of the Red Hat Enterprise Linux 6.9 Release Notes.		
Revision 0.0-4	Thu Jan 05 2017	Lenka Špačková
Release of the Red Hat Enterprise Linux 6.9 Beta Release Notes.		