



Red Hat Enterprise Linux 7

7.6 Release Notes

Release Notes for Red Hat Enterprise Linux 7.6

Red Hat Enterprise Linux 7 7.6 Release Notes

Release Notes for Red Hat Enterprise Linux 7.6

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

Legal Notice

Copyright © 2018 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 7.6 and document known problems in this release, as well as notable bug fixes, Technology Previews, deprecated functionality, and other details.

Table of Contents

PREFACE	11
CHAPTER 1. OVERVIEW	12
Security	12
Networking	12
Identity Management and Access Control	12
Management and Automation	12
Red Hat Insights	13
Red Hat Customer Portal Labs	13
CHAPTER 2. ARCHITECTURES	14
CHAPTER 3. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	15
KERNEL PARAMETERS	15
NEW AND UPDATED /PROC/SYS/KERNEL/ ENTRIES	16
NEW /PROC/SYS/NET/CORE ENTRIES	17
PART I. NEW FEATURES	18
CHAPTER 4. GENERAL UPDATES	19
In-place upgrade from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7	19
CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY	20
Certificate System now supports additional strong ciphers by default	20
samba rebased to version 4.8.3	20
Directory Server rebased to version 1.3.8.4	21
Certificate System rebased to version 10.5.9	21
jss rebased to version 4.4.4	21
The CRMFPopClient utility supports CRMF requests without key archival	21
Certificate System automatically applies ECC profiles when setting up root CA with ECC certificates	21
Certificate System now adds the SAN extension to server certificates	21
A low-level API to create X.509 certificates and CRLs has been added to JSS	21
The pcsc-lite-ccid driver now has support for new smart card readers	21
The pam_pkcs11 module now has support for certificate chains	21
dnsssec-keymgr automates DNSSEC key rollovers	22
DNSSEC validation can be disabled for selected domains	22
SSSD on an IdM client can now authenticate against a specific AD site or AD DC	22
CHAPTER 6. CLUSTERING	23
Pacemaker now supports path, mount, and timer systemd unit files	23
Support for Red Hat Enterprise Linux High Availability clusters on Alibaba Cloud	23
Support for Red Hat Enterprise Linux high availability clusters on Google Compute Cloud	23
New volume_group_check_only parameter for lvm resource agent	23
Support for VDO resource agent	23
The pcs command now supports filtering resource failures by an operation and its interval	23
New pcs commands to list available watchdog devices and test watchdog devices	23
CHAPTER 7. COMPILER AND TOOLS	24
The Net::SMTP Perl module now supports SSL	24
The Net::LDAP Perl module no longer defaults to TLS 1.0	24
timemaster now supports bonding devices	24
pcp rebased to version 4.1.0	24
The ps utility now displays the Login ID associated with processes	25
gcc-libraries rebased to version 8.2.1	25

systemtap rebased to version 3.3	25
GDB can disassemble instructions for the z14 processor of IBM Z architecture	25
New packages: java-11-openjdk	25
Support for new locales in glibc	26
New OFD Locking constants for 64-bit-offset programs	26
CHAPTER 8. DESKTOP	27
The sane-backends package is now built with systemd support	27
FreeType rebased to version 2.8	27
Nvidia Volta-based graphics cards are now supported	27
xorg-x11-server rebased to version 1.20.0-0.1	27
CHAPTER 9. FILE SYSTEMS	28
The CephFS kernel client is fully supported with Red Hat Ceph Storage 3	28
XFS now supports modifying labels on mounted file systems	28
pNFS SCSI layout is now fully supported for client and server	28
ima-evm-utils is now fully supported on AMD64 and Intel 64	28
CHAPTER 10. HARDWARE ENABLEMENT	29
genwqe-tools rebased to version 4.0.20 on IBM POWER	29
CHAPTER 11. INSTALLATION AND BOOTING	30
A new network-scripts option: IFDOWN_ON_SHUTDOWN	30
Improved content of error messages in network-scripts	30
Bootting from an iSCSI device that is not configured using iBFT is now supported	30
Installing and booting from NVDIMM devices is now supported	30
The --noghost option has been added to the rpm -V command	30
CHAPTER 12. KERNEL	31
The kdump FCoE target has been added into the kexec-tools documents	31
The SCHED_DEADLINE scheduler class enabled	31
User mount namespaces now fully supported	31
kernel.shmmax and kernel.shmall updated to kernel defaults on IBM Z	31
Updated aQuantia Corporation atlantic Network driver	31
Thunderbolt 3 is now supported	31
Intel® Omni-Path Architecture (OPA) Host Software	31
opal-prd rebased to version 6.0.4 on the little-endian variant of IBM POWER Systems	31
The SEV feature has been introduced for AMD virtual machines	32
CHAPTER 13. REAL-TIME KERNEL	33
About Red Hat Enterprise Linux for Real Time Kernel	33
kernel-rt sources updated	33
The SCHED_DEADLINE scheduler class for real time kernel fully supported	33
rt-entsk prevents IPI generation and delay of realtime tasks	33
CHAPTER 14. NETWORKING	34
Support for the libnftnl and nftables packages	34
ECMP fib_multipath_hash_policy support added to the kernel for IPv4 packets	34
Support for hardware time stamping on VLAN interfaces	34
Support for specifying speed and duplex 802-3-ethernet properties when 802-3-ethernet.auto-negotiation is enabled	34
Support for changing the DUID for IPv6 DHCP connections	34
ipset rebased to Linux kernel version 4.17	34
ipset (userspace) rebased to version 6.38	35
firewalld rebased to version 0.5.3	35

The ipset comment extension is now supported	35
radvd rebased to version 2.17	35
The default version for SMB now is auto-negotiated to the highest supported versions, SMB2 or SMB3	35
position in an nftables add or insert rule is replaced by handle and index	35
New features in net-snmp	35
firewalld-cmd --check-config now checks the validity of XML configuration files	36
CHAPTER 15. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE	37
Selected roles of Red Hat Enterprise Linux System Roles are now fully supported	37
CHAPTER 16. SECURITY	38
Clevis now supports TPM 2.0	38
gnutls rebased to 3.3.29	38
AES-GCM operations with OpenSSL are now faster on IBM z14	38
sudo rebased to version 1.8.23	38
usbguard rebased to version 0.7.4	39
audit rebased to 2.8.4	39
RPM now provides audit events	39
SELinux now supports extended_socket_class	39
selinux-policy now checks file permissions when mmap() is used	39
The RHEL7 DISA STIG profile now matches STIG Version 1, Release 4	40
Libreswan now supports PKCS #7-formatted X.509 certificates	40
libreswan rebased to version 3.25	40
openssl-ibmca rebased to version 2.0.0	40
sudo now runs PAM stack even when no authentication is required	40
cvsudoers converts between different sudoers formats	40
SCAP Security Guide now supports OSPP v4.2	40
selinux-policy now contains five additional SELinux booleans	40
CHAPTER 17. SERVERS AND SERVICES	42
rear rebased to version 2.4	42
The rear package now includes a user guide	42
The pcsc-lite interface now supports up to 32 devices	43
tuned rebased to version 2.10.0	43
The STOU FTP command has improved algorithm for generating unique file names	43
rsyslog imfile now supports symlinks	43
New rsyslog module: omkafka	43
New rsyslog module: mmkubernetes	43
CHAPTER 18. STORAGE	44
NVMe driver rebased to version 4.17-rc1	44
NVMe/FC is fully supported on Broadcom Emulex Fibre Channel Adapters	44
DM Multipath now enables blacklisting or whitelisting paths by protocol	44
New %0 wildcard added for the multipathd show paths format command to show path failures	44
New all_tg_pt multipath configuration option	45
CHAPTER 19. SYSTEM AND SUBSCRIPTION MANAGEMENT	46
cockpit rebased to version 173	46
reposync now by default skips packages whose location falls outside the destination directory	46
The yum clean all command now prints a disk usage summary	46
The yum versionlock plug-in now displays which packages are blocked when running the yum update command	46
The repotrack command now supports the --repofrompath option	47
Subscription manager now respects proxy_port settings from rhsm.conf	47

New package: sos-collector	47
CHAPTER 20. VIRTUALIZATION	48
virt-v2v converts virtual machine CPU topology	48
virt-v2v can import virtual machines directly to RHV	48
The i6300esb watchdog is now supported by libvirt	48
Paravirtualized clock added to Red Hat Enterprise Linux VMs	48
VNC console is supported on IBM Z	48
QEMU Guest Agent diagnostics enhanced	48
CHAPTER 21. ATOMIC HOST AND CONTAINERS	49
Red Hat Enterprise Linux Atomic Host	49
CHAPTER 22. RED HAT SOFTWARE COLLECTIONS	50
PART II. NOTABLE BUG FIXES	51
CHAPTER 23. AUTHENTICATION AND INTEROPERABILITY	52
Directory Server now supports certificates with all ciphers supported by NSS	52
Directory Server correctly generates the CSN	52
The client-cert-request utility no longer fails to create CSRs for ECC certificates	52
The pkiconsole utility no longer accepts ACLs with an empty expression	52
CMC CRMF requests using ECC keys work correctly	52
Installing Certificate System subsystems with ECC keys no longer fail	52
Directory Server clients are no longer randomly restricted by anonymous resource limits	52
Thread processing in Directory Server has been serialized	53
Deleting the memberOf attribute in Directory Server works correctly	53
The PBKDF2_SHA256 password storage scheme can now be used in Directory Server	53
Directory Server no longer crashes when removing connections from an active list	53
The Disk Monitoring feature shuts down Directory Server on low disk space	53
Directory Server no longer logs a warning when searching a non-existent DN in entrydn attributes	53
The pwdhash utility no longer crashes when using the CRYPT password storage scheme	54
The Directory Server Pass-through plug-in now supports encrypted connections using the STARTTLS command	54
Using the password policy feature works correctly if chain on update is enabled	54
Improved performance when the fine-grained password policy is enabled in Directory Server	54
Directory Server now retrieves members of the replica bind DN group when the first session is started	54
Creating a Directory Server back end with the name default is now supported	54
Updated Directory Server SNMP MIB definitions	54
rpc.yppasswdd now updates passwords also with SELinux disabled	55
The default of the nsslapd-enable-nunc-stans parameter has been changed to off	55
CHAPTER 24. CLUSTERING	56
PCS is able to find a token and connect to a node with upper case characters in its node name	56
pcs now shows correct value for failcount	56
At cluster startup, corosync starts on each node with a small delay to reduce the risk of JOIN flood	56
New /etc/sysconfig/pcsd option to reject client-initiated SSL/TLS renegotiation	56
CHAPTER 25. COMPILER AND TOOLS	57
GDB registers unaligned watchpoint hits on the 64-bit ARM architecture	57
Retpoline support in GCC on IBM Z architecture	57
binutils linker no longer terminates unexpectedly when encountering relocations against absolute address	57
The helper to store credentials in a GNOME keyring is now available in the git-gnome-keyring subpackage	57
git instaweb now works without any additional configuration and it is available in a separate subpackage	57
The man utility no longer prints gimme gimme gimme after midnight	57

sysctl now allows tuned to reset kernel parameters	58
llvm-private no longer crashes when used together with more recent libstdc++ library versions	58
ncat now correctly sets environment variables in UDP mode	58
ncat no longer uses the default HTTP port for all proxy types	58
Decoding and conversion of JPEG 2000 images now work correctly	58
strip no longer malforms binary files built with tools that use a later BFD library version	58
CHAPTER 26. HARDWARE ENABLEMENT	59
The lsslot -cpci command now correctly reports PCI slot types	59
The drmgr -C command now loads the rpadlpar_io kernel module	59
Diagnostic utilities now display CPU frequency values correctly	59
The ppc64_cpu utility no longer fails when reading CPU frequency	59
CHAPTER 27. INSTALLATION AND BOOTING	60
The network service no longer hangs on stop or restart	60
KSH no longer fails to process /etc/init.d/functions	60
Diskless NFS clients no longer hang when unmounting the root file system	60
A non-functioning systemctl reload network.service has been removed	60
Text mode will now prompt for a passphrase if a Kickstart file does not provide one while enabling encryption	60
A cmdline Kickstart installation with conflicting packages now displays an error message	60
The custom partitioning screen now displays relevant storage configuration error messages	60
Host name is now configured correctly on an installed system	61
The repart Kickstart command will now only create partitions that are required by the hardware platform	61
Installation started with boot option zfcpl.allow_lun_scan is applied to the installed system	61
The clearpart Kickstart command can now be used on disk partitions	61
CHAPTER 28. KERNEL	62
libcgroup no longer truncates the values of cgroup subsystem parameters that are longer than 100 characters	62
The mlx5 device no longer contains a firmware issue	62
CHAPTER 29. REAL-TIME KERNEL	63
A race condition that prevented tasks from being scheduled properly has been fixed	63
CHAPTER 30. NETWORKING	64
Bad offload warnings are no longer displayed using virtio_net	64
The L2TP sequence number handling now works correctly	64
The kernel no longer crashes when a tunnel_key mode is not specified	64
The sysctl net.ipv4.route.min_pmtu setting no longer set invalid values	64
wpa_supplicant no longer responds to packets whose destination address does not match the interface address	64
NetworkManager no longer fails to detect duplicate IPv4 addresses	64
firewalld now prevents partially applied rules	64
The wpa_supplicant upgrade no longer causes disconnections	65
CHAPTER 31. SECURITY	66
CardOS 5.3 smart cards with ECDSA support work correctly in OpenSC	66
Non-CCID-compliant smart card readers work in OpenSC	66
The pkcs11-tool utility now supports mechanism IDs and handles ECDSA keys correctly	66
OpenSCAP RPM verification rules no longer work incorrectly with VM and container file systems	66
sudo no longer blocks poll() for /dev/ptmx	66
CHAPTER 32. SERVERS AND SERVICES	67
pxlcolor and pxlmono now work correctly	67
The nuxwdog service starts correctly when a sub-CA is installed	67
Augeas reads /etc/fstab with white spaces more reliably	67

CHAPTER 33. STORAGE	68
mpathpersist no longer fails when opening too many files	68
The multipathd readsector0 checker now returns the correct result	68
DM Multipath is much less likely to output an incorrect timeout error	68
multipath now correctly prints the sysfs state of paths	68
multipathd can now correctly set APTPL when registering keys on path devices	68
CHAPTER 34. SYSTEM AND SUBSCRIPTION MANAGEMENT	69
The yum updateinfo commands now respect skip_if_unavailable option	69
PART III. TECHNOLOGY PREVIEWS	70
CHAPTER 35. GENERAL UPDATES	71
The systemd-importd VM and container image import and export service	71
CHAPTER 36. AUTHENTICATION AND INTEROPERABILITY	72
Use of AD and LDAP sudo providers	72
DNSSEC available as Technology Preview in IdM	72
Identity Management JSON-RPC API available as Technology Preview	72
The Custodia secrets service provider is now available	72
Containerized Identity Management server available as Technology Preview	73
CHAPTER 37. CLUSTERING	74
The pcs tool now manages bundle resources in Pacemaker	74
New fence-agents-heuristics-ping fence agent	74
Heuristics supported in corosync-qdevice as a Technology Preview	74
New LVM and LVM lock manager resource agents	74
CHAPTER 38. DESKTOP	76
Wayland available as a Technology Preview	76
Fractional Scaling available as a Technology Preview	76
CHAPTER 39. FILE SYSTEMS	77
ext4 and XFS file systems now support DAX	77
pNFS block layout is now available	77
OverlayFS	77
Btrfs file system	78
ima-evm-utils available as a Technology Preview for certain architectures	78
CHAPTER 40. HARDWARE ENABLEMENT	79
LSI Syncro CS HA-DAS adapters	79
tss2 enables TPM 2.0 for IBM Power LE	79
The ibmvnic device driver available as a Technology Preview	79
CHAPTER 41. INSTALLATION AND BOOTING	80
Custom system image creation with Composer available as a Technology Preview	80
CHAPTER 42. KERNEL	81
Heterogeneous memory management included as a Technology Preview	81
criu rebased to version 3.5	81
kexec as a Technology Preview	81
kexec fast reboot as a Technology Preview	81
perf cqm has been replaced by resctrl	81
TC HW offloading available as a Technology Preview	82
AMD xgbe network driver available as a Technology Preview	82

CHAPTER 43. NETWORKING	83
Cisco usNIC driver	83
Cisco VIC kernel driver	83
Trusted Network Connect	83
SR-IOV functionality in the qlcnict driver	83
The flower classifier with off-loading support	83
CHAPTER 44. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE	84
The postfix role of Red Hat Enterprise Linux System Roles as a Technology Preview	84
CHAPTER 45. SECURITY	85
USBGuard enables blocking USB devices while the screen is locked as a Technology Preview	85
pk12util can now import certificates signed with RSA-PSS	85
Support for certificates signed with RSA-PSS in certutil has been improved	85
NSS is now able to verify RSA-PSS signatures on certificates	85
SECCOMP can be now enabled in libreswan	85
CHAPTER 46. STORAGE	87
Multi-queue I/O scheduling for SCSI	87
Targetd plug-in from the libStorageMgmt API	87
Support for Data Integrity Field/Data Integrity Extension (DIF/DIX)	87
SCSI-MQ as a Technology Preview in the qla2xxx and lpfc drivers	87
NVMe/FC available as a Technology Preview in Qlogic adapters using the qla2xxx driver	87
CHAPTER 47. SYSTEM AND SUBSCRIPTION MANAGEMENT	89
YUM 4 available as Technology Preview	89
CHAPTER 48. VIRTUALIZATION	90
eBPF system call for tracing	90
USB 3.0 support for KVM guests	90
Select Intel network adapters now support SR-IOV as a guest on Hyper-V	90
No-IOMMU mode for VFIO drivers	90
virt-v2v can now use vmx configuration files to convert VMware guests	90
virt-v2v can convert Debian and Ubuntu guests	90
Virtio devices can now use vIOMMU	90
virt-v2v converts VMWare guests faster and more reliably	91
Open Virtual Machine Firmware	91
GPU-based mediated devices now support the VNC console	91
PART IV. DEVICE DRIVERS	92
CHAPTER 49. NEW DRIVERS	93
Network Drivers	93
Storage Drivers	93
Graphics Drivers and Miscellaneous Drivers	93
CHAPTER 50. UPDATED DRIVERS	94
Storage Driver Updates	94
Network Driver Updates	94
Graphics Driver and Miscellaneous Driver Updates	95
CHAPTER 51. DEPRECATED FUNCTIONALITY	96
Python 2 has been deprecated	96
LVM libraries and LVM Python bindings have been deprecated	96
Mirrored mirror log has been deprecated in LVM	96
Deprecated packages related to Identity Management and security	96

The Clevis HTTP pin has been deprecated	98
3DES is removed from the Python SSL default cipher list	98
sssd-secrets has been deprecated	98
Support for earlier IdM servers and for IdM replicas at domain level 0 will be limited	98
Bug-fix only support for the nss-pam-ldapd and NIS packages in the next major release of Red Hat Enterprise Linux	98
Use the Go Toolset instead of golang	99
mesa-private-llvm will be replaced with llvm-private	99
libdbi and libdbi-drivers have been deprecated	99
Ansible deprecated in the Extras channel	99
signtool has been deprecated	100
TLS compression support has been removed from nss	100
Public web CAs are no longer trusted for code signing by default	100
Sendmail has been deprecated	100
dmraid has been deprecated	100
Automatic loading of DCCP modules through socket layer is now disabled by default	100
rsyslog-libdbi has been deprecated	100
The inputname option of the rsyslog imudp module has been deprecated	100
SMBv1 is no longer installed with Microsoft Windows 10 and 2016 (updates 1709 and later)	101
FedFS has been deprecated	101
Btrfs has been deprecated	101
tcp_wrappers deprecated	101
nautilus-open-terminal replaced with gnome-terminal-nautilus	101
sslwrap() removed from Python	101
Symbols from libraries linked as dependencies no longer resolved by ld	101
Windows guest virtual machine support limited	102
libnetlink is deprecated	102
S3 and S4 power management states for KVM have been deprecated	102
The Certificate Server plug-in udnPwDirAuth is discontinued	102
Red Hat Access plug-in for IdM is discontinued	102
The Ipsilon identity provider service for federated single sign-on	102
Several rsyslog options deprecated	102
Deprecated symbols from the memkind library	102
Options of Sockets API Extensions for SCTP (RFC 6458) deprecated	103
Managing NetApp ONTAP using SSLv2 and SSLv3 is no longer supported by libstorageMgmt	103
dconf-dbus-1 has been deprecated and dconf-editor is now delivered separately	103
FreeRADIUS no longer accepts Auth-Type := System	104
Deprecated Device Drivers	104
Deprecated Adapters	107
The libcxgb3 library and the cxgb3 firmware package have been deprecated	112
SFN4XXX adapters have been deprecated	112
Software-initiated-only FCoE storage technologies have been deprecated	112
Containers using the libvirt-lxc tooling have been deprecated	113
The Perl and shell scripts for Directory Server have been deprecated	113
The gnome-shell-browser-plugin subpackage has been deprecated	113
The VDO read cache has been deprecated	113
cpuid has been deprecated	113
KDE has been deprecated	114
The lwresd daemon has been deprecated	114
PART V. KNOWN ISSUES	115
CHAPTER 52. AUTHENTICATION AND INTEROPERABILITY	116

RADIUS proxy functionality is now also available in IdM running in FIPS mode	116
CHAPTER 53. COMPILER AND TOOLS	117
GCC thread sanitizer included in RHEL no longer works	117
CHAPTER 54. DESKTOP	118
Firefox 60.1 ESR fails to start on IBM Z and POWER	118
GV100GL graphics cannot use correctly more than one monitor	118
The Files application can not burn disks in default installation	118
The on screen keyboard feature not visible in GTK applications	118
32- and 64-bit fwupd packages cannot be used together when installing or upgrading the system	118
Installation in and booting into graphical mode are not possible on Huawei servers	118
X.org server crashes during fast user switching	119
X.org X11 crashes on Lenovo T580	119
Soft lock-ups might occur during boot in the kernel with i915	119
System boots to a blank screen when Xinerama is enabled	119
CHAPTER 55. FILE SYSTEMS	120
Mounting a non-existent NFS export outputs a different error than in RHEL 6	120
XFS disables per-inode DAX functionality	120
CHAPTER 56. INSTALLATION AND BOOTING	121
Certain RPM packages are not available on binary DVDs	121
The content location detection code is not working on Red Hat Virtualization Hosts	121
Composer can not create live ISO system images	121
CHAPTER 57. KERNEL	122
Cache information is missing in sysfs if firmware does not support ACPI PPTT	122
PCI-passthrough of devices connected to PCIe slots is not possible with default settings of HPE ProLiant Gen8 and Gen9	122
Attaching a non-RoCE device to RXE driver no longer causes a kernel to panic	122
Enabling the BCC packages for the 64-bit AMD and Intel architectures only	122
Kernel panics can occur on virtual machines that use SEV	122
Branch prediction of ternary operators no longer causes a system panic	122
RAID1 write-behind causes a kernel panic	123
CHAPTER 58. NETWORKING	124
Verification of signatures using the MD5 hash algorithm is disabled in Red Hat Enterprise Linux 7	124
CHAPTER 59. SECURITY	125
OpenSCAP rpmverifypackage does not work correctly	125
dconf databases are not checked by OVAL	125
SCAP Workbench fails to generate results-based remediations from tailored profiles	125
OpenSCAP scanner results contain a lot of SELinux context error messages	125
CHAPTER 60. SERVERS AND SERVICES	126
Rsyslog cannot proceed if the default maximum of open files is exceeded	126
Tuned does not set kernel boot command line parameters	126
CHAPTER 61. STORAGE	127
LVM does not support event-based autoactivation of incomplete volume groups	127
The vdo service is disabled after upgrading to Red Hat Enterprise Linux 7.6	127
Data corruption occurs on RAID 10 reshape on top of VDO.	127
System boot is sometimes delayed by ndctl	127
LVM might cause data corruption in the first 128kB of allocatable space of a physical volume	127

CHAPTER 62. SYSTEM AND SUBSCRIPTION MANAGEMENT **129**

 Red Hat Satellite 5.8 availability of RHEL 7.6 EUS, AUS, TUS, and E4S streams delayed 129

APPENDIX A. COMPONENT VERSIONS **130**

APPENDIX B. LIST OF BUGZILLAS BY COMPONENT **131**

APPENDIX C. REVISION HISTORY **139**

PREFACE

Red Hat Enterprise Linux (RHEL) minor releases are an aggregation of individual security, enhancement, and bug fix errata. The *Red Hat Enterprise Linux 7.6 Release Notes* document describes the major changes made to the Red Hat Enterprise Linux 7 operating system and its accompanying applications for this minor release, as well as known problems and a complete list of all currently available Technology Previews.

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

Packages distributed with this release are listed in [Red Hat Enterprise Linux 7 Package Manifest](#). Migration from Red Hat Enterprise Linux 6 is documented in the [Migration Planning Guide](#).

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

CHAPTER 1. OVERVIEW

Security

- Driven by Trusted Platform Module (TPM) 2.0 hardware modules, the **Policy-Based Decryption (PBD)** capability has been extended to provide two layers of security for hybrid-cloud operations: the network-based mechanism is applicable in the cloud, while the use of TPM on-premises helps to keep information on disks physically more secure.
- The **GnuTLS** library now provides improved Hardware Security Module (HSM) support.
- **OpenSSL** now works with new CP Assist for Cryptographic Functions (CPACF) instructions to accelerate Galois/Counter Mode (GCM) of operation as available with IBM z14.
- Red Hat Certificate System distributed with Red Hat Enterprise Linux 7.6 provides new default cryptographic algorithms for RSA and ECC, which help maintain FIPS compliance and stay current with cryptography requirements from NIST and other standards bodies, as well as organizations responsible for handling sensitive information.

See [Chapter 16, Security](#) and [Chapter 5, Authentication and Interoperability](#) for more information.

Networking

- For better integration with counter-intrusion measures, firewall operations through Red Hat Enterprise Linux have been improved with enhancements to **nftables**. The **nft** command-line tool can now also provide improved control packet filtering, providing better overall visibility and simplified configuration for systems security.

For details, see [Chapter 14, Networking](#).

Identity Management and Access Control

- This release of OpenSC supports support new smart cards, for example, models with CardOS 5.3.

For details, see [Chapter 31, Security](#).

Management and Automation

- The tools for managing Red Hat Enterprise Linux 7 continue to be refined, with the latest version introducing enhancements to the Red Hat Enterprise Linux Web Console including:
 - Showing available updates on the system summary page
 - Automatic configuration of single sign-on for identity management, helping to simplify this task for security administrators
 - An interface to control firewall services
- The following Red Hat Enterprise Linux System Roles are now fully supported: **selinux**, **kdump**, **network**, and **timesync**.
- The integration of the **Extended Berkeley Packet Filter (eBPF)** provides a safer, more efficient mechanism for monitoring activity within the kernel and will help to enable additional performance monitoring and network tracing tools in the future. The **eBPF** tool is available as a Technology Preview.

For detailed information, refer to [Chapter 19, System and Subscription Management](#), [Chapter 15, Red Hat Enterprise Linux System Roles Powered by Ansible](#) and [Chapter 48, Virtualization](#).

Red Hat Insights

Since Red Hat Enterprise Linux 7.2, 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](#) or through Red Hat Satellite. To register your systems, follow the [Getting Started Guide for Insights](#).

Red Hat Customer Portal Labs

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

- [Registration Assistant](#)
- [Red Hat Code Browser](#)
- [Kickstart Generator](#)
- [Spectre And Meltdown Detector](#)
- [Load Balancer Configuration Tool](#)
- [L1 Terminal Fault Vulnerability Detector](#)
- [Red Hat Product Certificates](#)
- [Product Life Cycle Checker](#)
- [Red Hat Satellite Upgrade Helper](#)
- [Log Reaper](#)

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 7.6 is available on the following architectures: ^[1]

- 64-bit AMD
- 64-bit Intel
- IBM POWER7+ (big endian)
- IBM POWER8 (big endian) ^[2]
- IBM POWER8 (little endian) ^[3]
- IBM POWER9 (little endian) ^[4] ^[5]
- IBM Z ^[4] ^[6]
- 64-bit ARM ^[4]

^[1] Note that the Red Hat Enterprise Linux 7.6 installation is supported only on 64-bit hardware. Red Hat Enterprise Linux 7.6 is able to run 32-bit operating systems, including previous versions of Red Hat Enterprise Linux, as virtual machines.

^[2] Red Hat Enterprise Linux 7.6 POWER8 (big endian) are currently supported as KVM guests on Red Hat Enterprise Linux 7.6 POWER8 systems that run the KVM hypervisor.

^[3] Red Hat Enterprise Linux 7.6 POWER8 (little endian) is currently supported as a KVM guest on Red Hat Enterprise Linux 7.6 POWER8 systems that run the KVM hypervisor. In addition, Red Hat Enterprise Linux 7.6 POWER8 (little endian) guests are supported on Red Hat Enterprise Linux 7.6 POWER9 systems that run the KVM hypervisor in POWER8-compatibility mode on version 4.14 kernel using the kernel-alt package.

^[4] This architecture is supported with the kernel version 4.14, provided by the kernel-alt packages. For details, see the [Red Hat Enterprise Linux 7.5](#).

^[5] Red Hat Enterprise Linux 7.6 POWER9 (little endian) is currently supported as a KVM guest on Red Hat Enterprise Linux 7.6 POWER9 systems that run the KVM hypervisor on version 4.14 kernel using the kernel-alt package.

^[6] Red Hat Enterprise Linux 7.6 for IBM Z (both the 3.10 kernel version and the 4.14 kernel version) is currently supported as a KVM guest on Red Hat Enterprise Linux 7.6 for z Systems hosts that run the KVM on version 4.14 kernel using the kernel-alt package.

CHAPTER 3. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS

This chapter provides system administrators with a summary of significant changes in the kernel shipped with Red Hat Enterprise Linux 7.6. These changes include added or updated **proc** entries, **sysctl**, and **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes.

KERNEL PARAMETERS

hardened_usercopy = [KNL]

This parameter specifies whether hardening is enabled (default) or not enabled for the boot.

Hardened usercopy checking is used to protect the kernel from reading or writing beyond known memory allocation boundaries as a proactive defense against bounds-checking flaws in the kernel's **copy_to_user()**/**copy_from_user()** interface.

The valid settings are: **on**, **off**.

on – Perform hardened usercopy checks (default).

off – Disable hardened usercopy checks.

no-vmw-sched-clock [X86,PV_OPS]

Disables paravirtualized VMware scheduler clock and uses the default one.

rdt = [HW,X86,RDT]

Turns on or off individual RDT features.

Available features are: **cmt**, **mbmtotal**, **mbmlocal**, **l3cat**, **l3cdp**, **l2cat**, **l2cdp**, **mba**.

For example, to turn on **cmt** and turn off **mba**, use:

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

nospec_store_bypass_disable [HW]

Disables all mitigations for the Speculative Store Bypass vulnerability.

For more in-depth information about the Speculative Store Bypass (SSB) vulnerability, see [Kernel Side-Channel Attack using Speculative Store Bypass - CVE-2018-3639](#).

spec_store_bypass_disable = [HW]

Certain CPUs are vulnerable to an exploit against a common industry wide performance optimization known as Speculative Store Bypass.

In such cases, recent stores to the same memory location cannot always be observed by later loads during speculative execution. However, such stores are unlikely and thus they can be detected prior to instruction retirement at the end of a particular speculation execution window.

In vulnerable processors, the speculatively forwarded store can be used in a cache side channel attack. An example of this is reading memory to which the attacker does not directly have access, for example inside the sandboxed code.

This parameter controls whether the Speculative Store Bypass (SSB) optimization to mitigate the SSB vulnerability is used.

Possible values are:

on – Unconditionally disable SSB.

off – Unconditionally enable SSB.

auto – Kernel detects whether the CPU model contains an implementation of SSB and selects the most appropriate mitigation.

prctl – Controls SSB for a thread using prctl. SSB is enabled for a process by default. The state of the control is inherited on fork.

Not specifying this option is equivalent to **spec_store_bypass_disable=auto**.

For more in-depth information about the Speculative Store Bypass (SSB) vulnerability, see [Kernel Side-Channel Attack using Speculative Store Bypass - CVE-2018-3639](#).

nmi_watchdog = [KNL,BUGS=X86]

These settings can now be accessed at runtime with the use of the **nmi_watchdog** and **hardlockup_panic** sysctls.

NEW AND UPDATED /PROC/SYS/KERNEL/ ENTRIES

hardlockup_panic

This parameter controls whether the kernel panics if a hard lockup is detected.

Possible values are:

0 – Do not panic on hard lockup.

1 – Panic on hard lockup.

This can also be set using the **nmi_watchdog** kernel parameter.

perf_event_mlock_kb

Controls size of per-cpu ring buffer not counted against mlock limit.

The default value is **512 + 1** page.

perf_event_paranoid

Controls use of the performance events system by unprivileged users (without **CAP_SYS_ADMIN**).

The default value is **2**.

Possible values are:

-1 – Allow use of the majority of events by all users.

>=0 – Disallow ftrace function tracepoint and raw tracepoint access by users without **CAP_SYS_ADMIN**.

>=1 – Disallow CPU event access by users without **CAP_SYS_ADMIN**.

>=2 – Disallow kernel profiling by users without **CAP_SYS_ADMIN**.

NEW /PROC/SYS/NET/CORE ENTRIES

bpf_jit_harden

Enables hardening for the Berkeley Packet Filter (BPF) Just in Time (JIT) compiler.

Supported are Extended Berkeley Packet Filter (eBPF) JIT backends. Enabling hardening trades off performance, but can mitigate JIT spraying.

Possible values are:

0 – Disable JIT hardening (default value).

1 – Enable JIT hardening for unprivileged users only.

2 – Enable JIT hardening for all users.

PART I. NEW FEATURES

This part documents new features and major enhancements introduced in Red Hat Enterprise Linux 7.6.

CHAPTER 4. 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

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Migration_Planning_Guide/chap-Red_Hat_Enterprise_Linux-Migration_Planning_Guide-Upgrading.html and <https://access.redhat.com/solutions/637583>.

Note that the **Preupgrade Assistant** and the **Red Hat Upgrade Tool** are available in the Red Hat Enterprise Linux 6 Extras channel, see <https://access.redhat.com/support/policy/updates/extras>. (BZ#1432080)

CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY

Certificate System now supports additional strong ciphers by default

With this update, the following additional ciphers, which are compliant with the Federal Information Processing Standard (FIPS), are enabled by default in Certificate System:

- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_RSA_WITH_AES_256_GCM_SHA384

For a full list of enabled ciphers, enter:

```
# /usr/lib64/nss/unsupported-tools/listsuites | grep -B1 --no-group-separator "Enabled"
```

If you use a Hardware Security Module (HSM) with Certificate System, see the documentation of the HSM for supported ciphers. (BZ#[1550786](#))

samba rebased to version 4.8.3

The samba packages have been upgraded to upstream version 4.8.3, which provides a number of bug fixes and enhancements over the previous version:

- The **smbd** service no longer queries user and group information from Active Directory domain controllers and NT4 primary domain controllers directly. Installations with the **security** parameter set to **ads** or **domain** now require that the **winbindd** service is running.
- The dependency on global lists of trusted domains within the **winbindd** process has been reduced. For installations that do not require the global list, set the **winbind scan trusted domains** parameter in the **/etc/samba/smb.conf** file to **no**. For more information, see the parameter's description in the **smb.conf(5)** man page.
- The trust properties displayed in the output of the **wbinfo -m --verbose** command have been changed to correctly reflect the status of the system where the command is executed.
- Authentication from users of a one-way trust now works correctly when using the **idmap_rid** and **idmap_autorid** ID mapping back ends.

Samba automatically updates its tdb database files when the **smbd**, **nmbd**, or **winbind** daemon starts. Back up the databases files before starting Samba. Note that Red Hat does not support downgrading tdb database files.

For more information about notable changes, read the upstream release notes before updating:

<https://www.samba.org/samba/history/samba-4.8.0.html>. (BZ#1558560)

Directory Server rebased to version 1.3.8.4

The 389-ds-base packages have been upgraded to upstream version 1.3.8.4, which provides a number of bug fixes and enhancements over the previous version. For a complete list of notable changes, read the upstream release notes before updating:

- <http://directory.fedoraproject.org/docs/389ds/releases/release-1-3-8-4.html> (BZ#1560653)

Certificate System rebased to version 10.5.9

The pki-core packages have been upgraded to upstream version 10.5.9, which provides a number of bug fixes and enhancements over the previous version. (BZ#1557569)

jss rebased to version 4.4.4

The jss packages has been upgraded to upstream version 4.4.4, which provides a number of bug fixes and enhancements over the previous version. (BZ#1557575)

The CRMFPopClient utility supports CRMF requests without key archival

With this enhancement, users can create Certificate Request Message Format (CRMF) requests without the key archival option when using the CRMFPopClient utility. This feature increases flexibility because a Key Recovery Authority (KRA) certificate is no longer required. Previously, if the user did not pass the **-b transport_certificate_file** option to CRMFPopClient, the utility automatically used the KRA transport certificate stored in the transport.txt file. With this update, if **-b transport_certificate_file** is not specified, Certificate System creates a request without using key archival. (BZ#1585866)

Certificate System automatically applies ECC profiles when setting up root CA with ECC certificates

This update enhances Certificate System to automatically apply ECC profiles when setting up a new root CA with ECC profiles using the **pkispawn** utility. As a result, administrators no longer have to set the profile overwrite parameters for ECC certificates as a workaround in the configuration file passed to **pkispawn** when setting up a root CA. (BZ#1550742)

Certificate System now adds the SAN extension to server certificates

With this update, Certificate System adds the Subject Alternative Name (SAN) extension by default to server certificates and sets it to the Common Name (CN) of the certificate. (BZ#1562423)

A low-level API to create X.509 certificates and CRLs has been added to JSS

This enhancements adds a low-level API, which can be used to create X.509 certificate and certificate revocation lists (CRL) to the Java Security Services (JSS). (BZ#1560682)

The pcsc-lite-ccid driver now has support for new smart card readers

Previously, the **pcsc-lite-ccid** driver did not detect certain smart card readers. This enhancement adds the USB-ID values of these readers to the driver. As a result, **pcsc-lite-ccid** now detects the smart card readers in the described scenario.

Note that Red Hat did not test the smart card readers whose USB-ID have been added. (BZ#1558258)

The pam_pkcs11 module now has support for certificate chains

This update enhances the **pam_pkcs11** module to support Public Key Infrastructure for X.509 (PKIX) certificate chains. This enables more complex chain processing, including multiple paths to the leaf certificate. As a result, **pam_pkcs11** now validates PKIX certificate chains. (BZ#1578029)

dnssec-keymgr automates DNSSEC key rollovers

This update introduces **dnssec-keymgr**, a utility to automate DNS Security Extensions (DNSSEC) key rollovers. **dnssec-keymgr** enables automatic long-term management of DNS keys for secure zones due to its simple configurable policy. This makes it possible to roll out keys seamlessly, without interrupting the DNS service. (BZ#1510008)

DNSSEC validation can be disabled for selected domains

Previously, if DNSSEC validation was enabled and a specific domain was failing, no hosts in that domain could be reached. With this release, you can configure exemptions from DNS Security Extensions (DNSSEC) validation for selected zones if the validation fails because of incorrect configuration, not an attack. The addresses of the hosts in the failing domain are resolved as unsigned and can be reached, while all other names are validated for security risks. (BZ#1452091)

sssd on an IdM client can now authenticate against a specific AD site or AD DC

The **System Security Services Daemon** (SSSD) running on an Identity Management (IdM) client in a domain with a trust relationship with Active Directory (AD) can now be pinned to authenticate against a configured AD site or a configured set of AD Domain Controllers (DC).

Previously, **SSSD** relied completely on DNS SRV discovery done by **libkrb5**. However, this did not take AD sites into account because **libkrb5** has no notion of AD sites. If the administrator wanted to pin **SSSD** to authenticate against a set of AD DCs, they had to set the correct Key Distribution Centre (KDC) in the **/etc/krb5.conf** file, which was non-intuitive.

The enhancement is especially convenient for large environments, in which modifying the **/etc/krb5.conf** file on each client individually was previously the only available solution. (BZ#1416528)

CHAPTER 6. CLUSTERING

Pacemaker now supports `path`, `mount`, and `timer` systemd unit files

Previously, Pacemaker supported `service` and `socket` systemd unit files, but any other unit file type would be treated as a `service` unit and fail. With this release, `path`, `mount`, and `timer` systemd units can now be managed by a Pacemaker cluster. (BZ#1590483)

Support for Red Hat Enterprise Linux High Availability clusters on Alibaba Cloud

Red Hat Enterprise Linux 7.6 supports High Availability clusters of virtual machines (VMs) on Alibaba Cloud (Aliyun). For information on configuring a Red Hat Enterprise Linux High Availability Cluster on Alibaba Cloud, see <https://access.redhat.com/articles/3467251>. (BZ#1568589)

Support for Red Hat Enterprise Linux high availability clusters on Google Compute Cloud

Red Hat Enterprise Linux 7.6 supports high availability clusters of virtual machines (VMs) on Google Compute Cloud (GCP). For information on configuring a Red Hat Enterprise Linux High Availability Cluster on GCP, see <https://access.redhat.com/articles/3479821>. (BZ#1568588)

New `volume_group_check_only` parameter for `lvm` resource agent

The `lvm` resource agent now supports the `volume_group_check_only` parameter. When this parameter is set, only the volume group is checked when running a monitoring operation. Setting this parameter can be used to avoid timeouts with tagged volumes.

WARNING: This parameter should be used only when you have issues with timeouts, and when you must use the `lvm` resource agent and not the `LVM-activate` agent. (BZ#1470840)

Support for VDO resource agent

Red Hat Enterprise Linux now provides support for the `vdo-vol` resource agent to manage VDO (Virtual Data Optimizer) volumes as a high availability resource. (BZ#1538689)

The `pcs` command now supports filtering resource failures by an operation and its interval

Pacemaker now tracks resource failures per a resource operation on top of a resource name, and a node. The `pcs resource failcount show` command now allows filtering failures by a resource, node, operation, and interval. It provides an option to display failures aggregated per a resource and node or detailed per a resource, node, operation, and its interval. Additionally, the `pcs resource failcount reset` command now allows filtering failures by a resource, node, operation, and interval. (BZ#1427273)

New `pcs` commands to list available watchdog devices and test watchdog devices

In order to configure SBD with Pacemaker, a functioning watchdog device is required. The Red Hat Enterprise Linux 7.6 release supports the `pcs stonith sbd watchdog list` command to list available watchdog devices on the local node, and the `pcs stonith sbd watchdog test` command to test a watchdog device. (BZ#1475318)

CHAPTER 7. COMPILER AND TOOLS

The `Net::SMTP` Perl module now supports SSL

This update adds support for implicit and explicit TLS and SSL encryption to the `Net::SMTP` Perl module. As a result, it is now possible to communicate with SMTP servers through a secured channel. (BZ#[1557574](#))

The `Net::LDAP` Perl module no longer defaults to TLS 1.0

Previously, when the `Net::LDAP` Perl module was used for upgrading an unsecured LDAP connection to a TLS-protected one, the module used the TLS protocol version 1.0, which is currently considered insecure. With this update, the default TLS version has been removed from `Net::LDAP`, and both implicit (LDAPS schema) and explicit (LDAP schema) TLS protocols rely on the default TLS version selected in the `IO::Socket::SSL` Perl module. As a result, it is no longer necessary to override the TLS version in the `Net::LDAP` clients by passing the `sslversion` argument to the `start_tls()` method to preserve security. (BZ#[1520364](#))

`timemaster` now supports bonding devices

The `timemaster` program can be used to synchronize the system clock to all available time sources in case that there are multiple PTP domains available on the network, or fallback to NTP is needed.

This update adds the possibility to specify bonding devices in the active-backup mode in the `timemaster` configuration file. `timemaster` now checks if the active interface supports software or hardware timestamping and starts `ptp4l` on the bonding interface. (BZ#[1549015](#))

`pcp` rebased to version 4.1.0

The `pcp` packages have been upgraded to upstream version of Performace Co-Pilot 4.1.0, which provides a number of bug fixes and enhancements over the previous version:

- Added a sized-based interim compression to the `pmlogger_check(1)` script to reduce data volume sizes on systems configured via the `pcp-zeroconf` package.
- Daily compressed archive metadata files.
- Changed metric labels to first class PCP metric metadata.
- Metric help text and labels are now stored in PCP archives.
- Added more Linux kernel metrics: virtual machines, TTYs, aggregate interrupt and softirq counters, `af_unix/udp/tcp` connection (`inet/ipv6`), VFS locking, login sessions, AIO, capacity per block device, and other.
- Performance Metrics Application Programming Interface (PMAPI) and the Performance Metrics Domain Agent (PMDA) API have been refactored, including promotion and deprecation of individual functions.
- Added new virtual data optimizer (VDO) metrics to `pmdadm(1)`.
- Improved integration with Zabbix agentd service with further low-level-discovery support in the `pcp2zabbix(1)` function.
- Added a new PMDA `pmdabcc(1)` for exporting BCC and eBPF trace instrumentation.
- Added a new PMDA `pmdaprometheus(1)` to consume metrics from Prometheus end-points. (BZ#[1565370](#))

The **ps** utility now displays the Login ID associated with processes

The new format option **luid** of the **ps** utility now enables you to display the Login ID associated with processes.

To display the login ID attributes of running processes, use the following command:

```
$ ps -o luid
```

(BZ#[1518986](#))

gcc-libraries rebased to version 8.2.1

The gcc-libraries packages have been updated to upstream version 8.2.1. This update adds the following changes:

- The **libgfortran.so.5** and **libgfortran.so.4** Fortran libraries have been added to enable running applications built with Red Hat Developer Toolset versions 7 and later.
- The **libquadmath** library has been added as a dependency of the **libgfortran.so.5** library.
- The Cilk+ library has been removed. (BZ#1600265)

systemtap rebased to version 3.3

The systemtap packages have been upgraded to upstream version 3.3, which provides a number of bug fixes and enhancements:

- Limited support for the extended Berkeley Packet Filter (eBPF) tracing on the Intel64 and AMD64 architectures has been added. Use the **--runtime=bpf** option to use eBPF as a backend. Due to numerous limitations of **eBPF** and its SystemTap interface, only simple scripts work. For more information, see the Knowledge article <https://access.redhat.com/articles/3550581> and the `stapbpf(8)` manual page.
- The **--sysroot** option has been optimized for cross-compiled environments.
- A new **--example** option allows you to search the example scripts distributed with SystemTap without providing the whole path of the file.
- The SystemTap runtime and tapsets are compatible with kernel versions up to 4.17.
- Usage of SystemTap on systems with real time kernel or machines with a high number of CPUs has been improved.
- Handling of code used for Spectre and Meltdown attack mitigation has been improved. (BZ#[1565773](#))

GDB can disassemble instructions for the z14 processor of IBM Z architecture

The **GDB** debugger has been extended to disassemble instructions of the z14 processor of the IBM Z architecture, including guarded storage instructions. Previously, **GDB** displayed only the numerical values of such instructions in the **.long 0xNNNN** form. With this update, **GDB** can correctly display mnemonic names of assembly instructions in code targeting this processor. (BZ#1553104)

New packages: java-11-openjdk

The java-11-openjdk packages provide **OpenJDK 11** support through the **yum** utility.

OpenJDK 11 is the next Long-Term Support (LTS) version of Java supported by Red Hat after **OpenJDK**

8. It provides multiple new features including Modularization, Application Class Data Sharing, Heap Allocation on Alternative Memory Devices, Local-Variable Syntax for Lambda Parameters, and TLS 1.3 support.

The java-11-openjdk packages do not include unversioned **provides** because **OpenJDK 11** is not fully compatible with **OpenJDK 8**. (BZ#1570856)

Support for new locales in glibc

This update adds support for two new locales: Urdu (ur_IN) and Wolaytta (wal_ET). Additional support has also been added for newer currency symbols like the Euro, such as in **e1_GR@euro**. Users can now specify these locales using the relevant environment variables to take advantage of the new localization support. (BZ#[1448107](#))

New OFD Locking constants for 64-bit-offset programs

Open File Descriptor (OFD) locks are superior to per-process locks for some applications. With this update, 64-bit-offset programs (those that have **#define _FILE_OFFSET_BITS 64**) are able to use the **F_OFD_*** constants in system calls, although they still need to detect if the kernel supports those operations. Note that programs which use 32-bit file offsets do not have access to these constants, as the RHEL 7 ABI does not support translating them. (BZ#[1461231](#))

CHAPTER 8. DESKTOP

The sane-backends package is now built with systemd support

Scanner Access Now Easy (SANE) is a universal scanner interface whose backend's and library's features are provided by the sane-backends package. This update brings the following changes to SANE:

- The sane-backends package is built with systemd support.
- The saned daemon can be run without the need to create unit files manually, because these files are now shipped with sane-backends. (BZ#1512252)

FreeType rebased to version 2.8

The **FreeType** font engine has been rebased to version 2.8, which is required by GNOME 3.28. The 2.8 version has been modified to be API and Application Binary Interface (ABI) compatible with the previous version 2.4.11. (BZ#1576504)

Nvidia Volta-based graphics cards are now supported

This update adds support for Nvidia Volta-based graphics cards. As a result, the **modesetting** user-space driver, which is able to handle the basic operations and single graphic output, is used. However, 3D graphic is handled by the **llvmpipe** driver because Nvidia did not share public signed firmware for 3D. To reach maximum performance of the card, use the Nvidia binary driver. (BZ#1457161)

xorg-x11-server rebased to version 1.20.0-0.1

The xorg-x11-server packages have been rebased to upstream version 1.20.0-0.1, which provides a number of bug fixes and enhancements over the previous version:

- Added support for the following input devices: Wacom Cintiq Pro 24, Wacom Cintiq Pro 32 tablet, Wacom Pro Pen 3D.
- Added support for Intel Cannon Lake and Whiskey Lake platform GPUs.
- Added support for S3TC texture compression in OpenGL
- Added support for X11 backing store **always** mode.
- Added support for Nvidia Volta series of graphics.
- Added support for AMD Vega graphics and Raven APU. (BZ#1564632)

CHAPTER 9. FILE SYSTEMS

The CephFS kernel client is fully supported with Red Hat Ceph Storage 3

The Ceph File System (CephFS) kernel module enables Red Hat Enterprise Linux nodes to mount Ceph File Systems from Red Hat Ceph Storage clusters. The kernel client in Red Hat Enterprise Linux is a more efficient alternative to the Filesystem in Userspace (FUSE) client included with Red Hat Ceph Storage. Note that the kernel client currently lacks support for CephFS quotas.

The CephFS kernel client was introduced in Red Hat Enterprise Linux 7.3 as a Technology Preview, and since the release of Red Hat Ceph Storage 3, CephFS is fully supported.

For more information, see the Ceph File System Guide for Red Hat Ceph Storage 3:

https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/3/html/ceph_file_system_guide/. (BZ#1205497)

XFS now supports modifying labels on mounted file systems

You can now modify the label attribute of mounted XFS file systems using the `xfs_io` utility:

```
# xfs_io -c "label new-label" /mount-point
```

Previously, it was only possible to modify labels on unmounted file systems using the `xfs_admin` utility, which is still supported. (BZ#1322930)

pNFS SCSI layout is now fully supported for client and server

Client and server support for parallel NFS (pNFS) SCSI layouts is now fully supported. It was first introduced in Red Hat Enterprise Linux 7.3 as a Technology Preview.

Building on the work of block layouts, the pNFS layout is defined across SCSI devices and contains sequential series of fixed-size blocks as logical units that must be capable of supporting SCSI persistent reservations. The Logical Unit (LU) devices are identified by their SCSI device identification, and fencing is handled through the assignment of reservations. (BZ#1305092)

ima-evm-utils is now fully supported on AMD64 and Intel 64

The `ima-evm-utils` package is now fully supported when used on the AMD64 and Intel 64 architecture. Note that on other architectures, `ima-evm-utils` remains in Technology Preview.

The `ima-evm-utils` package provides utilities to label the file system and verify the integrity of your system at run time using the Integrity Measurement Architecture (IMA) and Extended Verification Module (EVM) features. These utilities enable you to monitor if files have been accidentally or maliciously altered. (BZ#1627278)

CHAPTER 10. HARDWARE ENABLEMENT

genwqe-tools rebased to version 4.0.20 on IBM POWER

The genwqe-tools packages have been rebased to version 4.0.20 for IBM POWER architectures. This version provides a number of bug fixes and enhancements over the previous version, most notably:

- CompressBound has been fixed
- Debugging tools have been added
- The **genwqe_cksum** tool has been fixed
- Missing manual pages in the spec file have been fixed
- New compiler warnings have been fixed
- **Z_STREAM_END** detection circumvention has been improved (BZ#1521050)

CHAPTER 11. INSTALLATION AND BOOTING

A new network-scripts option: IFDOWN_ON_SHUTDOWN

This update adds the **IFDOWN_ON_SHUTDOWN** option for **network-scripts**. Setting this option to **yes**, **true**, or leaving it empty has no effect. If you set this option to **no**, or **false**, it causes the **ifdown** calls to not be issued when stopping or restarting the **network** service.

This can be useful in situations where NFS (or other network file system) mounts are in a stale state, because the network was shut down before the mount was cleanly unmounted. (BZ#[1583677](#))

Improved content of error messages in network-scripts

The network-scripts now display more verbose error messages when the installation of bonding drivers fails. (BZ#[1542514](#))

Bootting from an iSCSI device that is not configured using iBFT is now supported

This update provides a new installer boot option **inst.nonibftiscsiboot** that supports the installation of boot loader on an iSCSI device that has not been configured in the iSCSI Boot Firmware Table (iBFT).

This update helps in a use case where the iSCSI device is not configured in the iBFT for installation, it is added manually by using the **iscsi** Kickstart command or the installer GUI; the iBFT is not used for booting the installed system from the iSCSI device, an iPXE boot from SAN features is used. (BZ#[1562301](#))

Installing and booting from NVDIMM devices is now supported

Prior to this update, Nonvolatile Dual Inline Memory Module (NVDIMM) devices in any mode were ignored by the installer.

With this update, kernel improvements to support NVDIMM devices provide improved system performance capabilities and enhanced file system access for write-intensive applications like database or analytic workloads, as well as reduced CPU overhead.

This update introduces support for:

- The use of NVDIMM devices for installation using the **nvdimm** Kickstart command and the GUI, making it possible to install and boot from NVDIMM devices in sector mode and reconfigure NVDIMM devices into sector mode during installation.
- The extension of **Kickstart** scripts for **Anaconda** with commands for handling NVDIMM devices.
- The ability of **grub2**, **efibootmgr**, and **efivar** system components to handle and boot from NVDIMM devices. (BZ#[1612965](#), BZ#[1280500](#), BZ#[1590319](#), BZ#[1558942](#))

The --noghost option has been added to the rpm -v command

This update adds the **--noghost** option to the **rpm -V** command. If used with this option, **rpm -V** verifies only the non-ghost files that were altered, which helps diagnose system problems. (BZ#[1395818](#))

CHAPTER 12. KERNEL

The **kdump FCoE target** has been added into the **kexec-tools** documents

This update adds the **kdump** Fibre Channel over Ethernet (FCoE) target into the **kexec-tools** documents. As a result, users now have better understanding about the state and details of **kdump** on FCoE target support. (BZ#1352763)

The **SCHED_DEADLINE** scheduler class enabled

This update adds support for the **SCHED_DEADLINE** scheduler class for the Linux kernel. The scheduler enables predictable task scheduling based on application deadlines. **SCHED_DEADLINE** benefits periodic workloads by guaranteeing timing isolation, which is not based only on a fixed priority but also on the applications' timing requirements. (BZ#1344565)

User mount namespaces now fully supported

The mount namespaces feature, previously available as a Technology Preview, is now fully supported. (BZ#1350553)

kernel.shmmax and **kernel.shmall** updated to kernel defaults on IBM Z

Previously, applications that required a large amount of memory in some cases terminated unexpectedly due to low values of the **kernel.shmmax** and **kernel.shmall** parameters on IBM Z. This update aligns the values of **kernel.shmmax** and **kernel.shmall** with kernel defaults, which helps avoid the described crashes. (BZ#1493069)

Updated aQuantia Corporation **atlantic** Network driver

The aQuantia Corporation Network driver, **atlantic.ko.xz**, has been updated to version 2.0.2.1-kern and it is now fully supported. (BZ#1451438)

Thunderbolt 3 is now supported

This update adds support for the Thunderbolt 3 interface. (BZ#1620372)

Intel® Omni-Path Architecture (OPA) Host Software

Intel Omni-Path Architecture (OPA) host software is fully supported in Red Hat Enterprise Linux 7.6.

Intel OPA provides Host Fabric Interface (HFI) hardware with initialization and setup for high performance data transfers (high bandwidth, high message rate, low latency) between compute and I/O nodes in a clustered environment.

For instructions on installing Intel Omni-Path Architecture documentation, see:

https://www.intel.com/content/dam/support/us/en/documents/network-and-i-o/fabric-products/Intel_OP_Software_RHEL_7_6_RN_K34562.pdf (BZ#1627126)

opal-prd rebased to version 6.0.4 on the little-endian variant of IBM POWER Systems

On the little-endian variant of IBM POWER Systems, the **opal-prd** packages have been upgraded to upstream version 6.0.4, which provides a number of bug fixes and enhancements over the previous version. For example:

- Performance in High Performance Computing (HPC) environments has been improved.
- The **powernv_flash** module is now explicitly loaded on systems based on Baseboard Management Controller (BMC), which ensures that the flash device is created before the **opal-prd** daemon starts.

- Error on the first failure for soft or hard offline is no longer displayed by the **opa1-prd** daemon. (BZ#[1564097](#), BZ#1537001)

The SEV feature has been introduced for AMD virtual machines

The AMD Secure Encrypted Virtualization (SEV) is now available in Red Hat Enterprise Linux 7. SEV enables the memory contents of a virtual machine (VM) to be transparently encrypted with a key unique to the guest VM. The memory controller contains a high-performance encryption engine that can be programmed with multiple keys for use by a different VMs in the system. The programming and management of these keys is handled by the AMD Secure Processor firmware. As a result, SEV significantly increases the security of virtual machines that use the AMD architecture.

Note that the **swiotlb=262144** parameter currently has be to added to the guest kernel command line to ensure that virtual machines that use SEV are stable. For details, see the Known Issues section. (BZ#1361286)

CHAPTER 13. REAL-TIME KERNEL

About Red Hat Enterprise Linux for Real Time Kernel

The Red Hat Enterprise Linux for Real Time Kernel is designed to enable fine-tuning for systems with extremely high determinism requirements. The major increase in the consistency of results can, and should, be achieved by tuning the standard kernel. The real-time kernel enables gaining a small increase on top of increase achieved by tuning the standard kernel.

The real-time kernel is available in the **rhel-7-server-rt-rpms** repository. The [Installation Guide](#) contains the installation instructions and the rest of the documentation is available at [Product Documentation for Red Hat Enterprise Linux for Real Time](#).

kernel-rt sources updated

The kernel-rt sources have been upgraded to be based on the latest Red Hat Enterprise Linux kernel source tree, which provides a number of bug fixes and enhancements over the previous version. (BZ#[1553351](#))

The SCHED_DEADLINE scheduler class for real time kernel fully supported

The **SCHED_DEADLINE** scheduler class for the real-time kernel, which was introduced in Red Hat Enterprise Linux 7.4 as a Technology Preview, is now fully supported. The scheduler enables predictable task scheduling based on application deadlines. **SCHED_DEADLINE** benefits periodic workloads by guaranteeing timing isolation, which is based not only on a fixed priority but also on the applications' timing requirements. (BZ#[1297061](#))

rt-entsk prevents IPI generation and delay of realtime tasks

The chrony daemon, **chronyd**, enables or disables network timestamping, which activates a static key within the kernel. When a static key is enabled or disabled, three inter-processor interrupt (IPIs) are generated to notify other processors of the activation.

Previously, rapid activation and deactivation of the **chronyd** static keys led to a delay of a realtime task. Consequently, a latency spike occurred. With this update, systemd starts the **rt-entsk** program, which keeps timestamping enabled and prevents the IPIs from being generated. As a result, IPI generation no longer occurs in a rapid succession, and realtime tasks are no longer delayed due to this bug. (BZ#[1616038](#))

CHAPTER 14. NETWORKING

Support for the libnftnl and nftables packages

The nftables and libnftnl packages, previously available as a Technology Preview, are now supported.

The nftables packages provide a packet-filtering tool, with numerous improvements in convenience, features, and performance over previous packet-filtering tools. It is the designated successor to the **iptables**, **ip6tables**, **arptables**, and **ebtables** utilities.

The libnftnl packages provide a library for low-level interaction with nftables Netlink API over the **libmnl** library. (BZ#1332585)

ECMP fib_multipath_hash_policy support added to the kernel for IPv4 packets

This update adds support for Equal-cost multi-path routing (ECMP) hash policy choice using **fib_multipath_hash_policy**, a new **sysctl** setting that controls which hash policy to use for multipath routes. When **fib_multipath_hash_policy** is set to **1**, the kernel performs **L4 hash**, which is a multipath hash for IPv4 packets according to a **5-tuple** (source IP, source port, destination IP, destination port, IP protocol type) set of values. When **fib_multipath_hash_policy** is set to **0** (default), only **L3 hash** is used (the source and destination IP addresses).

Note that if you enable **fib_multipath_hash_policy**, the Internet Control Message Protocol (ICMP) error packets are not hashed according to the inner packet headers. This is a problem for anycast services as the ICMP packet can be delivered to the incorrect host. (BZ#1511351)

Support for hardware time stamping on VLAN interfaces

This update adds hardware time stamping on VLAN interfaces (driver dp83640 is excluded). This allows applications, such as **linuxptp**, to enable hardware time stamping. (BZ#1520356)

Support for specifying speed and duplex 802-3-ethernet properties when 802-3-ethernet.auto-negotiation is enabled

Previously, when **802-3-ethernet.auto-negotiation** was enabled on an Ethernet connection, all the **speed** and **duplex** modes supported by the Network Interface Card (NIC) were advertised. The only option to enforce a specific **speed** and **duplex** mode was to disable **802-3-ethernet.auto-negotiation** and set **802-3-ethernet.speed** and **802-3-ethernet.duplex** properties. This was not correct because the **1000BASE-T** and **10GBASE-T** Ethernet standards require **auto-negotiation** to be always enabled. With this update, you can enable a specific **speed** and **duplex** when **auto-negotiation** is enabled. (BZ#1487477)

Support for changing the DUID for IPv6 DHCP connections

With this update, users can configure the DHCP Unique Identifier (DUID) in **NetworkManager** to get an IPv6 address from a Dynamic Host Configuration Protocol (DHCP) server. As a result, users can now specify the DUID for DHCPv6 connections using the new property, **ipv6.dhcp-duid**. For more details on values set for **ipv6.dhcp-duid**, see the **nm-settings(5)** man page. (BZ#1414093)

ipset rebased to Linux kernel version 4.17

The **ipset** kernel component has been upgraded to upstream Linux kernel version 4.17 which provides a number of enhancements and bug fixes over the previous version. Notable changes include:

- The following **ipset** types are now supported:
- **hash:net,net**
- **hash:net,port,net**

- `hash:ip,mark`
- `hash:mac`
- `hash:ip,mac` (BZ#[1557599](#))

ipset (userspace) rebased to version 6.38

The `ipset (userspace)` package has been upgraded to upstream version 6.38, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- The userspace `ipset` is now aligned to the Red Hat Enterprise Linux (RHEL) kernel `ipset` implementation in terms of supported `ipset` types
- A new type of set, `hash:ipmac`, is now supported (BZ#[1557600](#))

firewalld rebased to version 0.5.3

The `firewalld` service daemon has been upgraded to upstream version 0.5.3, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Added the `--check-config` option to verify sanity of configuration files.
- Generated interfaces such as `docker0` are now correctly re-added to zones after `firewalld` restarts.
- A new IP set type, `hash:mac`, is now supported. (BZ#[1554993](#))

The ipset comment extension is now supported

This update adds the `ipset comment` extension. This enables you to add entries with a comment. For more information, see the `ipset (8)` man page. (BZ#[1496859](#))

radvd rebased to version 2.17

The **router advertisement daemon (radvd)** has been upgraded to version 2.17. The most notable change is that now `radvd` supports the selection of router advertisements source address. As a result, connection tracking no longer fails when the router's address is moved between hosts or firewalls. (BZ#[1475983](#))

The default version for SMB now is auto-negotiated to the highest supported versions, SMB2 or SMB3

With this update, the default version of the Server Message Block (SMB) protocol has been changed from SMB1 to be auto-negotiated to the highest supported versions SMB2 or SMB3. Users can still choose to explicitly mount with the less secure SMB1 dialect (for old servers) by adding the `vers=1.0` option on the Common Internet File System (CIFS) mount.

Note that SMB2 or SMB3 do not support Unix Extensions. Users that depend on Unix Extensions need to review the mount options and ensure that `vers=1.0` is used. (BZ#[1471950](#))

position in an nftables add or insert rule is replaced by handle and index

With this update of the `nftables` packages, the `position` parameter in an `add` or `insert` rule has been deprecated and replaced by the `handle` and `index` arguments. This syntax is more consistent with the `replace` and `delete` commands. (BZ#[1571968](#))

New features in net-snmp

The `net-snmp` package in Red Hat Enterprise Linux 7 has been extended with the following new features:

- net-snmp now supports monitoring disks of ZFS file system.
- net-snmp now supports monitoring disks of ASM Cluster (AC) file system. (BZ#[1533943](#), BZ#1564400)

firewalld-cmd --check-config now checks the validity of XML configuration files

This update introduces the **--check-config** option for the **firewall-cmd** and **firewall-offline-cmd** commands. The new option checks a user configuration of the **firewalld** daemon in XML files. The verification script reports syntax errors in custom rule definitions if any. (BZ#[1477771](#))

CHAPTER 15. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE

Selected roles of Red Hat Enterprise Linux System Roles are now fully supported

Red Hat Enterprise Linux System Roles provides a configuration interface for Red Hat Enterprise Linux subsystems, which makes system configuration easier through the inclusion of Ansible Roles. This interface enables managing system configurations across multiple versions of Red Hat Enterprise Linux, as well as adopting new major releases. The interface currently consists of the following roles:

- **selinux**
- **kdump**
- **network**
- **timesync**
- **postfix**

Red Hat Enterprise Linux System Roles have been available since Red Hat Enterprise Linux 7.4 as a Technology Preview. With this update, the **selinux**, **kdump**, **network**, and **timesync** roles are fully supported. The **postfix** role continues to be available as a Technology Preview.

Since Red Hat Enterprise Linux 7.4, the Red Hat Enterprise Linux System Roles packages have been distributed through the Extras channel. For details regarding Red Hat Enterprise Linux System Roles, see <https://access.redhat.com/articles/3050101>.

Selected roles of the `rhel-system-roles` package have received multiple bug fixes and significant enhancements to improve interface consistency, usability, and conformance to Ansible best practices. Note that for the **timesync**, **kdump**, and **selinux** roles, the changes are not backward compatible and it is necessary to update playbooks that use them. For more information, see <https://access.redhat.com/articles/3561071>. (BZ#1479381)

CHAPTER 16. SECURITY

clevis now supports TPM 2.0

With this update, the **clevis** pluggable framework for Policy-Based Decryption (PBD) supports also clients that encrypt using a Trusted Platform Module 2.0 (TPM 2.0) chip. For more information and the list of possible configuration properties, see the **clevis-encrypt-tpm2(1)** man page.

Note that this feature is available only on systems with the 64-bit Intel or 64-bit AMD architecture. (BZ#1472435)

gnutls rebased to 3.3.29

The GNU Transport Layer Security (GnuTLS) library has been upgraded to upstream version 3.3.29, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Improved the PKCS#11 cryptographic token interface for hardware security modules (HSMs): added DSA support in **p11tool** and fixed key import in certain Atos HSMs.
- Improved counter-measures for the TLS Cipher Block Chaining (CBC) record padding. The previous counter-measures had certain issues and were insufficient when the attacker had access to the CPU cache and performed a chosen-plaintext attack (CPA).
- Disabled the legacy **HMAC-SHA384** cipher suites by default. (BZ#1561481)

AES-GCM operations with openssl are now faster on IBM z14

This update introduces support for additional acceleration of cryptographical operations with new CP Assist for Cryptographic Functions (CPACF) instructions available on IBM z14 systems. As a result, **AES-GCM** operations with the **OpenSSL** library are now executed faster on IBM z14 and later hardware. (BZ#1519396)

sudo rebased to version 1.8.23

The sudo packages have been upgraded to upstream version 1.8.23, which provides a number of bug fixes and enhancements over the previous version:

- The new **cvtsudoers** utility replaces both the **sudoers2ldif** script and the **visudo -x** functionality. It can read a file in either sudoers or LDIF format and produce JSON, LDIF, or sudoers output. It is also possible to filter the generated output file by user, group, or host name.
- The **always_query_group_plugin** option is now set explicitly in the default **/etc/sudoers** file. Users who upgrade from previous versions and want to retain the old group-querying behavior should ensure that this setting is in place after the upgrade.
- PAM account management modules are now run even when no password is required.
- The new **case_insensitive_user** and **case_insensitive_group** sudoers options enable to control whether **sudo** does case-sensitive matching of users and groups in **sudoers**. Case-insensitive matching is now the default.
- It is now an error to specify the **runas** user as an empty string on the command line. Previously, an empty **runas** user was treated the same as an unspecified **runas** user.
- I/O log files are now created with group **ID 0** by default unless the **io_log_user** or **io_log_group** options are set in **sudoers**.

- It is now possible to preserve bash shell functions in the environment where the **env_reset** **sudoers** setting is disabled by removing the ***=() *** pattern from the **env_delete** list. (BZ#1547974)

usbguard rebased to version 0.7.4

The usbguard packages have been rebased to upstream version 0.7.4. This version provides a number of bug fixes and enhancements over the previous version, most notably:

- The **usbguard-daemon** now exits with an error if it fails to open a logging file or an audit event file.
- The present device enumeration algorithm is now more reliable. Enumeration timeouts no longer cause the **usbguard-daemon** process to exit.
- The **usbguard watch** command now includes the **-e** option to run an executable for every received event. The event data is passed to the executable through environment variables. (BZ#1508878)

audit rebased to 2.8.4

The audit packages have been upgraded to upstream version 2.8.4, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Added support for dumping internal state. You can now run the **service auditd state** command to see information about the **Audit** daemon.
- Added support for the **SOFTWARE_UPDATE** event generated by the **rpm** and **yum** tools.
- Allowed unlimited retries during a remote logging startup. This helps to start even if the aggregating server is not running when a client is booted.
- Improved IPv6 remote logging. (BZ#1559032)

RPM now provides audit events

With this update, the **RPM Package Manager** (RPM) provides audit events. The information that a software package is installed or updated is important for system analysis with the Linux **Audit** system. **RPM** now creates a **SOFTWARE_UPDATE** audit event whenever a package is installed or upgraded by the **root** user. (BZ#1555326)

SELinux now supports `extended_socket_class`

This update introduces the **extended_socket_class** policy capability that enables a number of new SELinux object classes to support all of the known network socket address families. It also enables the use of separate security classes for Internet Control Message Protocol (ICMP) and Stream Control Transmission Protocol (SCTP) sockets, which were previously mapped to the **rawip_socket** class. (BZ#1564775, BZ#1427553)

selinux-policy now checks file permissions when `mmap()` is used

This release introduces a new permission check on the **mmap()** system call. The purpose of a separate map permission check on **mmap()** is to permit policy to prohibit memory mapping of specific files for which you need to ensure that every access is revalidated. This is useful for scenarios where you expect the files to be relabeled at run-time to reflect state changes, for example, in a cross-domain solution or an assured pipeline without data copying.

This functionality is enabled by default. Also, a new SELinux boolean, **domain_can_mmap_files**, has been added. If **domain_can_mmap_files** is enabled, every domain can use **mmap()** in every file, a character device or a block device. If **domain_can_mmap_files** is disabled, the list of domains that

can use `mmap()` is limited. (BZ#1460322)

The RHEL7 DISA STIG profile now matches STIG Version 1, Release 4

With this update of the **SCAP Security Guide** project, the RHEL7 Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) profile is aligned with STIG Version 1, Release 4. Note that certain rules do not contain an automated check or fix. (BZ#1443551)

Libreswan now supports PKCS #7-formatted X.509 certificates

With this update, the **Libreswan** Virtual Private Network application supports also PKCS #7-formatted X.509 certificates. This enables interoperability with systems running Microsoft Windows. (BZ#1536404)

libreswan rebased to version 3.25

The libreswan packages have been upgraded to upstream version 3.25, which provides a number of bug fixes and enhancements over the previous version.

Note that previously, an incorrect configuration forbidding Perfect Forward Secrecy with the `pfs=no` option and setting an ESP/AH PFS `modp` group (for example, `esp=aes-sha2;modp2048`) would load and ignore the `modp` setting. With this update, these connections fail to load with the **ESP DH algorithm MODP2048 is invalid as PFS policy is disabled** error message. (BZ#1591817)

openssl-ibmca rebased to version 2.0.0

The openssl-ibmca packages have been upgraded to upstream version 2.0.0, which provides a number of bug fixes and enhancements over the previous version:

- The Elliptic-Curve Cryptography (ECC) functionality is now supported.
- Compatibility with various **OpenSSL** versions has been increased.

Note that to use the ECC functionality with a shared CEX4C adapter in the z/VM 6.4 system, the Authorized Program Analysis Report (APAR) VM65942 is required. (BZ#1519395)

sudo now runs PAM stack even when no authentication is required

With this update, the **sudo** utility runs Pluggable Authentication Module (PAM) account management modules even when the **NOPASSWD** option is configured in the policy. This enables checking for restrictions imposed by PAM modules outside of the authentication phase. As a result, PAM modules, such as **pam_time**, now work properly in the described scenario. (BZ#1533964)

cvtsudoers converts between different sudoers formats

The new **cvtsudoers** utility enables the administrator to convert rules between different **sudoers** security policy file formats. See the **cvtsudoers(1)** man page for the list of available options and examples of usage. (BZ#1548380)

SCAP Security Guide now supports OSPP v4.2

This update of the scap-security-guide packages introduces a new profile defining the core requirements of OSPP (General-Purpose Operating System Protection Profile) v4.2. The new profile ID is **ospp42**, and the previously released profile USGCB (United States Government Configuration Baseline) OSPP v4.0 is available with ID **ospp**. (BZ#1619689)

selinux-policy now contains five additional SELinux booleans

This update of the selinux-policy packages introduces the following SELinux booleans:

- **keepalived_connect_any** - allows the **keepalived** service to connect to arbitrary ports.

- **tomcat_use_execmem** - allows the **Tomcat** server to make its stack executable.
- **tomcat_can_network_connect_db** - allows **Tomcat** to connect to the **PosgreSQL** port.
- **redis_enable_notify** - allows the **redis-sentinel** service to run notification scripts.
- **zabbix_run_sudo** - allows the **zabbix_agent** service to run the **sudo** utility. (BZ#[1443473](#), BZ#[1565226](#), BZ#[1477948](#), BZ#1421326, BZ#[1347052](#))

CHAPTER 17. SERVERS AND SERVICES

rear rebased to version 2.4

The rear packages that provide the **Relax-and-Recover** tool (ReaR) have been upgraded to upstream version 2.4, which provides a number of bug fixes and enhancements over the previous version. Notably:

- The default behavior when resizing partitions in migration mode has been changed. Only the size of the last partition is now changed by default; the start positions of every partition are preserved. If the previous behavior is needed, set the **AUTORESIZE_PARTITIONS** configuration variable to **yes**. See the description of the configuration variables **AUTORESIZE_PARTITIONS**, **AUTORESIZE_EXCLUDE_PARTITIONS**, **AUTOSHRINK_DISK_SIZE_LIMIT_PERCENTAGE**, and **AUTOINCREASE_DISK_SIZE_THRESHOLD_PERCENTAGE** in the `/usr/share/rear/conf/default.conf` file for more information on how to control the partition resizing.
- The network setup now supports teaming (with the exception of Link Aggregation Control Protocol - LACP), bridges, bonding, and VLANs.
- Support for Tivoli Storage Manager (TSM) has been improved. In particular, support for the password store in the TSM client versions 8.1.2 and later has been added, fixing the bug where the generated ISO image did not support restoring the OS if those TSM versions were used for backup.
- Support for partition names containing blank and slash characters has been fixed.
- SSH secrets (private keys) are no longer copied to the recovery system, which prevents their leaking. As a consequence, SSH in the recovery system cannot use the secret keys from the original system. See the description of the **SSH_FILES**, **SSH_ROOT_PASSWORD**, and **SSH_UNPROTECTED_PRIVATE_KEYS** variables in the `/usr/share/rear/conf/default.conf` file for more information on controlling this behavior.
- Numerous improvements to support of the IBM POWER Systems architecture have been added, such as support for including the backup in the rescue ISO image and for multiple ISOs.
- Multipath support has been enhanced. For example, support for software RAID on multipath devices has been added.
- Support for secure boot has been added. The **SECURE_BOOT_BOOTLOADER** variable can be used for specifying any custom-signed boot loader.
- Support for restoring disk layout of software RAID devices with missing components has been added.
- The standard error and standard output channels of programs invoked by **ReaR** are redirected to the log file instead of appearing on the terminal. Programs prompting for user input on the standard output or standard error channel will not work correctly. Their standard output channel should be redirected to file descriptor **7** and standard input channel from file descriptor **6**. See the Coding Style documentation on the **ReaR** wiki for more details.
- Support for recovery of systems with LVM thin pool and thin volumes has been added. (BZ#1496518, BZ#1484051, BZ#1534646, BZ#1498828, BZ#1571266, BZ#1539063, BZ#1464353, BZ#1536023)

The rear package now includes a user guide

This update adds the user guide into the rear package, which provides the **Relax-and-Recover** tool (ReaR). After installation of rear, you can find the user guide in the `/usr/share/doc/rear-2.4/relax-and-recover-user-guide.html` file. (BZ#1418459)

The pcsc-lite interface now supports up to 32 devices

In Red Hat Enterprise Linux 7.6, the number of devices the **pcsc-lite** smart card interface supports has been increased from 16 to 32. (BZ#1516993)

tuned rebased to version 2.10.0

The tuned packages have been rebased to upstream version 2.10.0, which provides a number of bug fixes and enhancements over the previous version.

Notable changes include:

- an added mssql profile (shipped in a separate tuned-profiles-mssql subpackage)
- the **tuned-adm** tool now displays a relevant log snippet in case of error
- fixed verification of a CPU mask on systems with more than 32 cores (BZ#1546598)

The stou FTP command has improved algorithm for generating unique file names

The **STOU** FTP command allows transferring files to the server and storing them with unique file names. Previously, the **STOU** command created the names of the files by taking the file name, supplied as an argument to the command, and adding a numerical suffix and incrementing the suffix by one. In some cases, this led to a race condition. Subsequently the scripts which used **STOU** to upload files with the same file name could fail. This update modifies **STOU** to create unique file names in a way which helps to avoid the race condition and improves the functioning of scripts that use **STOU**. To enable the improved algorithm for generating unique file names using **STOU**, enable the **better_stou** option in the configuration file (usually `/etc/vsftpd/vsftpd.conf`) by adding the following line:

better_stou=YES (BZ#1479237)

rsyslog imfile now supports symlinks

With this update, the **rsyslog imfile** module delivers better performance and more configuration options. This enables to use the module for more complicated file monitoring use cases. Users of **rsyslog** are now able to use file monitors with glob patterns anywhere along the configured path and rotate symlink targets with increased data throughput when compared to the previous version. (BZ#1531295)

New rsyslog module: omkafka

To enable **kafka** centralized data storage scenarios, you can now forward logs to the **kafka** infrastructure using the new **omkafka** module. (BZ#1482819)

New rsyslog module: mmkubernetes

To enable scenarios using **rsyslog** in favor of other log collectors and where kubernetes container metadata are required, a new **mmkubernetes** module has been added to Red Hat Enterprise Linux. (BZ#1539193)

CHAPTER 18. STORAGE

NVMe driver rebased to version 4.17-rc1

The **NVMe** driver has been rebased to upstream version 4.17-rc1, which provides a number of bug fixes and enhancements over the previous version. Notable changes are as follows:

- added error handling improvements for Nonvolatile Memory Express (NVMe) over Remote Direct Memory Access (RDMA)
- added fixes for keeping connections over the RDMA transport alive

Note that the driver does not support the Data Integrity Field/Data Integrity Extension (DIF/DIX) Protection Information implementation, and does not support multipathing over NVMe-over-Fabrics transport. (BZ#1515584)

NVMe/FC is fully supported on Broadcom Emulex Fibre Channel Adapters

The NVMe over Fibre Channel (NVMe/FC) transport type is now fully supported in Initiator mode when used with Broadcom Emulex Fibre Channel 32Gbit adapters.

NVMe over Fibre Channel is an additional fabric transport type for the Nonvolatile Memory Express (NVMe) protocol, in addition to the Remote Direct Memory Access (RDMA) protocol that was previously introduced in Red Hat Enterprise Linux.

To enable NVMe/FC in the **lpfc** driver, edit the **/etc/modprobe.d/lpfc.conf** file and the following option:

```
lpfc_enable_fc4_type=3
```

This feature was introduced as a Technology Preview in Red Hat Enterprise Linux 7.5. Drivers other than **lpfc** still remain in Technology Preview. See the Technology Previews part for more information.

Additional restrictions:

- Multipath is not supported with NVMe/FC.
- NVMe clustering is not supported with NVMe/FC.
- Initiators do not support using NVMe/FC and SCSI/FC at the same time.
- The kernel-alt package does not support NVMe/FC.
- **kdump** is not supported with NVMe/FC.
- Booting from Storage Area Network (SAN) NVMe/FC is not supported. (BZ#1584753)

DM Multipath now enables blacklisting or whitelisting paths by protocol

Device Mapper Multipath (DM Multipath) now supports the **protocol** configuration option in the **blacklist** and **blacklist_exceptions** configuration sections. This enables you to blacklist or whitelist paths based on the protocol they use, such as **scsi** or **nvme**. For SCSI devices, you can also specify the transport: for example **scsi:fc** or **scsi:iscsi**. (BZ#1593459)

New **%0** wildcard added for the **multipathd show paths** format command to show path failures

The **multipathd show paths format** command now supports the **%0** wildcard to display path failures. Support for this wildcard makes it easier for users to track which paths have been failing in a multipath device. (BZ#[1554516](#))

New all_tg_pt multipath configuration option

The **defaults** and **devices** sections of the **multipath.conf** configuration file now support the **all_tg_pts** parameter, which defaults to **no**. If this option is set to **yes**, when **mpathpersist** registers keys it will treat a key registered from one host to one target port as going from one host to all target ports. Some arrays, notably the EMC VNX, treat reservations as between one host and all target ports. Without **mpathpersist** working the same way, it would give reservation conflicts. (BZ#1541116)

CHAPTER 19. SYSTEM AND SUBSCRIPTION MANAGEMENT

cockpit rebased to version 173

The cockpit packages, which provide the Cockpit browser-based administration console, have been upgraded to version 173. This version provides a number of bug fixes and enhancements. Notable changes include:

- The menu and navigation can now work with mobile browsers.
- **Cockpit** now supports alternate Kerberos keytabs for Cockpit's web server, which enables configuration of Single Sign-On (SSO).
- Automatic setup of Kerberos keytab for Cockpit web server.
- Automatic configuration of SSO with FreeIPA for **Cockpit** is possible.
- **Cockpit** requests FreeIPA SSL certificate for Cockpit's web server.
- **Cockpit** shows available package updates and missing registrations on system front page.
- A Firewall interface has been added.
- The flow control to avoid user interface hangs and unbounded memory usage for big file downloads has been added.
- Terminal issues in Chrome have been fixed.
- **Cockpit** now properly localizes numbers, times, and dates.
- Subscriptions page hang when accessing as a non-administrator user has been fixed.
- **Log in** is now localized properly.
- The check for root privilege availability has been improved to work for FreeIPA administrators as well. (BZ#[1568728](#), BZ#[1495543](#), BZ#[1442540](#), BZ#1541454, BZ#1574630)

reposync now by default skips packages whose location falls outside the destination directory

Previously, the **reposync** command did not sanitize paths to packages specified in a remote repository, which was insecure. A security fix for CVE-2018-10897 has changed the default behavior of **reposync** to not store any packages outside the specified destination directory. To restore the original insecure behavior, use the new **--allow-path-traversal** option. (BZ#1609302)

The yum clean all command now prints a disk usage summary

When using the **yum clean all** command, the following hint was always displayed:

```
Maybe you want: rm -rf /var/cache/yum
```

With this update, the hint has been removed, and **yum clean all** now prints a disk usage summary for remaining repositories that were not affected by **yum clean all** (BZ#[1481220](#))

The yum versionlock plug-in now displays which packages are blocked when running the yum update command

Previously, the **yum versionlock** plug-in, which is used to lock RPM packages, did not display any

information about packages excluded from the update. Consequently, users were not warned that such packages will not be updated when running the **yum update** command. With this update, **yum versionlock** has been changed. The plug-in now prints a message about how many package updates are being excluded. In addition, the new **status** subcommand has been added to the plug-in. The **yum versionlock status** command prints the list of available package updates blocked by the plug-in. (BZ#1497351)

The repotrack command now supports the --repofrompath option

The **--repofrompath** option, which is already supported by the **repoquery** and **repoclosure** commands, has been added to the **repotrack** command. As a result, non-root users can now add custom repositories to track without escalating their privileges. (BZ#1506205)

Subscription manager now respects proxy_port settings from rhsm.conf

Previously, subscription manager did not respect changes to the default **proxy_port** configuration from the **/etc/rhsm/rhsm.conf** file. Consequently, the default value of 3128 was used even after the user had changed the value of **proxy_port**.

With this update, the underlying source code has been fixed, and subscription manager now respects changes to the default **proxy_port** configuration. However, making any change to the **proxy_port** value in **/etc/rhsm/rhsm.conf** requires an selinux policy change. To avoid selinux denials when changing the default **proxy_port**, run this command for the benefit of the **rhsmcertd** daemon process:

```
semanage port -a -t squid_port_t -p tcp <new_proxy_port>
```

(BZ#1576423)

New package: sos-collector

sos-collector is a utility that gathers **sosreports** from multi-node environments. **sos-collector** facilitates data collection for support cases and it can be run from either a node or from an administrator's local workstation that has network access to the environment. (BZ#1481861)

CHAPTER 20. VIRTUALIZATION

virt-v2v converts virtual machine CPU topology

With this update, the **virt-v2v** utility preserves the CPU topology of the converted virtual machines (VMs). This ensures that the VM CPU works the same way after the conversion as it did before the conversion, which avoids potential runtime problems. (BZ#[1541908](#))

virt-v2v can import virtual machines directly to RHV

The **virt-v2v** utility is now able to output a converted virtual machine (VM) directly to a Red Hat Virtualization (RHV) client. As a result, importing VMs converted by **virt-v2v** using the Red Hat Virtualization Manager (RHVM) is now easier, faster, and more reliable.

Note that this feature requires RHV version 4.2 or later to work properly. (BZ#[1557273](#))

The i6300esb watchdog is now supported by libvirt

With this update, the **libvirt** API supports the i6300esb watchdog device. As a result, KVM virtual machines can use this device to automatically trigger a specified action, such as saving a core dump of the guest if the guest OS becomes unresponsive or terminates unexpectedly. (BZ#[1447169](#))

Paravirtualized clock added to Red Hat Enterprise Linux VMs

With this update, the paravirtualized **sched_clock()** function has been integrated in the Red Hat Enterprise Linux kernel. This improves the performance of Red Hat Enterprise Linux virtual machines (VMs) running on VMWare hypervisors.

Note that the function is enabled by default. To disable it, add the **no-vmw-sched-clock** option to the kernel command line. (BZ#[1507027](#))

VNC console is supported on IBM Z

This update enables the **virtio-gpu** kernel configuration in guests running on the IBM Z architecture. As a result, KVM guests on an IBM Z host are now able to use the VNC console to display their graphical output. (BZ#[1570090](#))

QEMU Guest Agent diagnostics enhanced

To maintain qemu-guest-agents compatibility with the latest version of VDSM, a number of features have been backported from the most recent upstream version.

These include the addition of **qemu-get-host-name**, **qemu-get-users**, **qemu-get-osinfo**, and **qemu-get-timezone** commands, which improve the diagnostic capabilities of QEMU Guest Agent. (BZ#[1569013](#))

CHAPTER 21. ATOMIC HOST AND CONTAINERS

Red Hat Enterprise Linux Atomic Host

Red Hat Enterprise Linux Atomic Host is a secure, lightweight, and minimal-footprint operating system optimized to run Linux containers. See the [Atomic Host and Containers Release Notes](#) for the latest new features, known issues, and Technology Previews.

CHAPTER 22. 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 7 on AMD64 and Intel 64 architectures, the 64-bit ARM architecture, IBM Z, and IBM POWER, little endian. Certain components are available also for all supported releases of Red Hat Enterprise Linux 6 on AMD64 and Intel 64 architectures.

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. Red Hat Developer Toolset is included 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. NOTABLE BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 7.6 that have a significant impact on users.

CHAPTER 23. AUTHENTICATION AND INTEROPERABILITY

Directory Server now supports certificates with all ciphers supported by NSS

Due to a restriction in Directory Server, administrators could only use RSA and Fortezza ciphers. As a consequence, certificates created with a different cipher, such as ECC certificates, were not supported. This update removes this restriction. As a result, administrators can now use certificates with all ciphers supported by the underlying Network Security Services (NSS) database when configuring TLS in Directory Server. (BZ#[1582747](#))

Directory Server correctly generates the CSN

In a Directory Server replication topology, updates are managed by using Change Sequence Numbers (CSN) based on time stamps. New CSNs must be higher than the highest CSN present in the replica update vector (RUV). In case the server generates a new CSN in the same second as the most recent CSN, the sequence number is increased to ensure that it is higher. However, if the most recent CSN and the new CSN were identical, the sequence number was not increased. In this situation, the new CSN was, except the replica ID, identical to the most recent one. As a consequence, a new update in the directory appeared in certain situations older than the most recent update. With this update, Directory Server increases the CSN if the sequence number is lower or equal to the most recent one. As a result, new updates are no longer considered older than the most recent data. (BZ#[1559945](#))

The `client-cert-request` utility no longer fails to create CSRs for ECC certificates

Previously, the `generatePkcs10Request` method in the Certificate System's `client-cert-request` utility failed to map the curve and length parameters. Consequently, the utility failed to create certificate signing requests (CSR) for Elliptic Curve Cryptography (ECC) certificates. The problem has been fixed. As a result, using `client-cert-request` for creating CSRs for ECC certificates works as expected. (BZ#[1549632](#))

The `pkiconsole` utility no longer accepts ACLs with an empty expression

The Certificate System server rejects saving invalid access control lists (ACL). As a consequence, when saving an ACL with an empty expression, the server rejected the update and the `pkiconsole` utility displayed an `StringIndexOutOfBoundsException` error. With this update, the utility rejects empty ACL expressions. As a result, invalid ACLs cannot be saved and the error is no longer displayed. (BZ#[1546708](#))

CMC CRMF requests using ECC keys work correctly

Previously, during verification, Certificate System encoded the ECC public key incorrectly in CMC Certificate Request Message Format (CRMF) requests. As a consequence, requesting an ECC certificate with Certificate Management over CMS (CMC) in CRMF failed. The problem has been fixed, and as a result, CMC CRMF requests using ECC keys work as expected. (BZ#[1580394](#))

Installing Certificate System subsystems with ECC keys no longer fail

Previously, due to a bug in the Certificate System installation procedure, installing a Key Recovery Authority (KRA) with ECC keys failed. To fix the problem, the installation process has been updated to handle both RSA and ECC subsystems automatically. As a result, installing subsystems with ECC keys no longer fail. (BZ#[1568615](#))

Directory Server clients are no longer randomly restricted by anonymous resource limits

Previously, Directory Server did not remember when the first operation, bind, or a connection was started. As a consequence, the server applied in certain situations anonymous resource limits to an authenticated client. With this update, Directory Server properly marks authenticated client connections. As a result, it applies the correct resource limits, and authenticated clients no longer get randomly restricted by anonymous resource limits. (BZ#[1515190](#))

Thread processing in Directory Server has been serialized

On an incoming replicated session, a replicated operation must only be processed when the previous one is completed. In certain situations, the thread which processed the start session operation continued to read and process replicated operations. Consequently, two replicated operations ran in parallel that led to inconsistencies, such as an completed child **add** operation before the parent entry was added. With this update, the thread processing the start session operation no longer processes further operations, even if some are available in the read buffer. As a result, the inconsistencies no longer occur in the mentioned scenario. (BZ#[1552698](#))

Deleting the **memberof** attribute in Directory Server works correctly

If an administrator moves a group in Directory Server from one subtree to another, the **memberof** plug-in deletes the **memberof** attribute with the old value and adds a new **memberof** attribute with the new group's distinguished name (DN) in affected user entries. Previously, if the old subtree was not within the scope of the **memberof** plug-in, deleting the old **memberof** attribute failed because the values did not exist. As a consequence, the plug-in did not add the new **memberof** value, and the user entry contained an incorrect **memberof** value. With this update, the plug-in now checks the return code when deleting the old value. If the return code is **no such value**, the plug-in only adds the new **memberof** value. As a result, the **memberof** attribute information is correct. (BZ#[1551071](#))

The PBKDF2_SHA256 password storage scheme can now be used in Directory Server

If a Red Hat Directory Server instance was installed using version 10.1.0 or earlier and subsequently updated, the update script did not enable the Password-Based Key Derivation Function version 2 (PBKDF2) plug-in. As a consequence, the **PBKDF2_SHA256** password storage scheme could not be used in the **nsslapd-rootpwstoragescheme** and **passwordStorageScheme** parameter. This update automatically enables the plug-in. As a result, administrators can now use the **PBKDF2_SHA256** password storage scheme. (BZ#[1576485](#))

Directory Server no longer crashes when removing connections from an active list

Directory Server manages established connections in an active list. When a thread flags a connection for closing, the server waits until there are no active threads left on the connection to remove the connection from the active list. In certain situations, the number of active threads is less than the actual number of threads. In this scenario, Directory Server moves the connection out of the active list and flags it as invalid. Another remaining thread which detects that the connection is invalid also attempts to remove it from the active list. However, the code that removes the connection from the active list expects that the connection has valid list pointers. If the pointers are invalid because the connection is not on the active list, Directory Server terminates unexpectedly. With this update, the server checks that the list pointers are valid before using them. As a result, the server no longer crashes when attempting to remove a connection from the active list. (BZ#[1566444](#))

The Disk Monitoring feature shuts down Directory Server on low disk space

Due to changes in the way Directory Server sets the error log level, the Disk Monitoring feature in Directory Server failed to detect that the error log level was set to the default level. As a consequence, Directory Server did not correctly shut down when the file system was full. The way the Disk Monitoring feature checks the error level has been updated. As a result, Disk Monitoring now correctly shuts down the server if the disk space is low. (BZ#[1568462](#))

Directory Server no longer logs a warning when searching a non-existent DN in **entrydn** attributes

Previously, searches for a non-existent distinguished name (DN) set in the **entrydn** attribute caused Directory Server to log a warning in the error log. With this update, the server correctly handles situations when an **entrydn** attribute fails to find a match. As a result, the server no longer logs a misleading warning. (BZ#[1570033](#))

The **pwdhash** utility no longer crashes when using the **CRYPT** password storage scheme

Previously, the **pwdhash** utility used an invalid mutex lock when creating a hash using the **CRYPT** password storage scheme. As a consequence, **pwdhash** failed with a segmentation fault error. With this update, the utility uses the re-entrant form of the **crypt()** function that does not require a lock. As a result, **pwdhash** no longer crashes when using the **CRYPT** password storage scheme. (BZ#[1570649](#))

The Directory Server **Pass-through** plug-in now supports encrypted connections using the **STARTTLS** command

Previously, the **Pass-through** plug-in in Directory Server did not support encrypted connections if the encryption was started using the **STARTTLS** command. The problem has been fixed, and the **Pass-through** plug-in now supports connections that use the **STARTTLS** command. (BZ#[1581737](#))

Using the password policy feature works correctly if **chain on update** is enabled

On a Directory Server read-only consumer, the **Password must be changed after reset** password policy setting was not enforced because the flag for marking the user that must change their password is set on the connection itself. If this setting was used with the **chain on update** feature, the flag was lost. As a consequence, the password policy feature did not work. With this update, the server sets the flag on **chain on update** connections properly. As a result, the password policy feature works correctly. (BZ#[1582092](#))

Improved performance when the fine-grained password policy is enabled in Directory Server

When a search evaluates the **shadowAccount** entry, Directory Server adds the shadow attributes to the entry. If the fine-grained password policy is enabled, the **shadowAccount** entry can contain its own **pwdpolicysubentry** policy attribute. Previously, to retrieve this attribute, the server started an internal search for each **shadowAccount** entry, which was unnecessary because the entry was already known to the server. With this update, Directory Server only starts internal searches if the entry is not known. As a result, the performance of searches, such as response time and throughput, is improved. (BZ#[1593807](#))

Directory Server now retrieves members of the replica bind DN group when the first session is started

Directory Server replicas define entries that are authorized to replicate updates to the replica itself. If the entries are members of the group set in the **nsds5replicabinddn** attribute, the group is retrieved periodically based on the interval set in the **nsDS5ReplicaBindDnGroupCheckInterval** attribute. If the entry is not a member at the time the server retrieves the group, any session that is authenticated using this entry is not authorized to replicate updates. This behavior remains until the entry becomes a member of the group and the server retrieves the group again. As a consequence, replication fails for the first interval set in **nsDS5ReplicaBindDnGroupCheckInterval**. With this update, the server retrieves the group when the first session is started rather than when the replica is created. As a result, the group is taken into account at the first attempt it is checked. (BZ#[1598478](#))

Creating a Directory Server back end with the name **default** is now supported

Previously, the name **default** was reserved in Directory Server. As a consequence, creating a back end named **default** failed. With this update, Directory Server no longer reserves this name, and administrators can create a back end named **default**. (BZ#[1598718](#))

Updated Directory Server SNMP MIB definitions

Previously, the Simple Network Management Protocol (SNMP) Management Information Base (MIB) definitions provided by the 389-ds-base package did not conform to the Structure of Management Information Version 2 (SMIv2) defined in RFC 2578. As a consequence, the **lint** utility reported errors.

The definitions have now been updated, and as a result, the MIB definitions comply with the SMlv2 specification (BZ#[1525256](#))

rpc.yppasswdd now updates passwords also with SELinux disabled

Previously, when the **SELinux** security module was disabled on the system, the **rpc.yppasswdd** update function failed to perform the update action. As a consequence, **rpc.yppasswdd** was unable to update the user password. With this update, **rpc.yppasswdd** checks whether **SELinux** is enabled on the system before detecting the **SELinux** context type for the **passwd** files. As a result, **rpc.yppasswdd** now correctly updates passwords in the described scenario. (BZ#[1492892](#))

The default of the nsslapd-enable-nunc-stans parameter has been changed to off

Previously, the nunc-stans framework was enabled by default in Directory Server, but the framework is not stable. As a consequence, deadlocks and file descriptor leaks could occur. This update changes the default value of the **nsslapd-enable-nunc-stans** parameter to **off**. As a result, Directory Server is now stable. (BZ#[1614501](#))

CHAPTER 24. CLUSTERING

PCS is able to find a token and connect to a node with upper case characters in its node name

Previously, PCS was unable to find a token for any node name with upper case characters, and it would report an error that the node is not authenticated. This occurred because the **pcs cluster auth** command would lowercase all node names before storing them to the PCS token file. With this fix, PCS does not lowercase node names before storing them to the PCS token file. (BZ#[1590533](#))

pcs now shows correct value for failcount

Starting with the Red Hat Enterprise Linux 7.5 release, the **pcs resource failcount show** command always showed a **failcount** of zero, even when this was not the correct value. This occurred because the format of resource failcounts was changed in Pacemaker. With this fix, the **pcs** utility is able to parse the new **failcount** format and it displays the correct value. (BZ#1588667)

At cluster startup, corosync starts on each node with a small delay to reduce the risk of JOIN flood

Starting **corosync** on all nodes at the same time may cause a JOIN flood, which may result in some nodes not joining the cluster. With this update, each node starts **corosync** with a small delay to reduce the risk of this happening. (BZ#[1572886](#))

New /etc/sysconfig/pcsd option to reject client-initiated SSL/TLS renegotiation

When TLS renegotiation is enabled on the server, a client is allowed to send a renegotiation request, which initiates a new handshake. Computational requirements of a handshake are higher on a server than on a client. This makes the server vulnerable to DoS attacks. With this fix, a new option has been added to the **/etc/sysconfig/pcsd** configuration file to reject renegotiations. Note that the client can still open multiple connections to a server with a handshake performed for all of the connections. (BZ#1566382)

CHAPTER 25. COMPILER AND TOOLS

GDB registers unaligned watchpoint hits on the 64-bit ARM architecture

Previously, the **GDB** debugger provided only limited support for unaligned hardware watchpoints used by the **watch**, **rwatch**, and **awatch** **GDB** commands on the 64-bit ARM architecture. As a consequence, **GDB** running on such systems failed to register some watchpoint hits and subsequently did not stop the debugged program.

GDB has been extended to handle this situation. As a result, it can correctly handle any hardware watchpoints on the 64-bit ARM architecture, including unaligned ones. (BZ#[1347993](#))

Retpoline support in GCC on IBM Z architecture

This update adds support for retpoline generation in the GNU Compiler Collection (GCC) on IBM Z architecture. Retpolines are a technique used by the kernel to reduce the overhead of mitigating Spectre Variant 2 attacks described in CVE-2017-5715. (BZ#[1552021](#))

binutils linker no longer terminates unexpectedly when encountering relocations against absolute address

Previously, the linker from the binutils package could not properly handle relocations against an absolute address. As a consequence, encountering such relocations caused a segmentation fault of the linker.

The linker has been extended to handle relocations against absolute addresses and the problem no longer occurs. (BZ#[1557346](#))

The helper to store credentials in a GNOME keyring is now available in the **git-gnome-keyring** subpackage

Previously, installing **git** automatically installed GNOME components as a dependency, because the helper to store credentials in a GNOME keyring was part of the **git** package. With this update, the helper has been moved into the separate **git-gnome-keyring** subpackage. As a result, the size of a **git** installation is reduced.

To install the subpackage:

```
# yum install git-gnome-keyring
```

(BZ#[1284081](#))

git instaweb now works without any additional configuration and it is available in a separate subpackage

Previously, the **git instaweb** command required a web server and did not work in the default installation. With this update, **git instaweb** has been moved into the separate **git-instaweb** subpackage, which depends on the **Apache** web server, and is configured to use the web server automatically. As a result, **git instaweb** now works without any further configuration when **git-instaweb** is installed. To install the subpackage:

```
# yum install git-instaweb
```

(BZ#[1213059](#))

The **man** utility no longer prints **gimme gimme gimme** after midnight

Prior to this update, there was an Easter egg in the **man** utility that printed **gimme gimme gimme** in the standard error output at 00:30 local time. As a consequence, under certain circumstances the unexpected output misled automated tools. With this update, the Easter egg has been removed, and the

described problem no longer occurs. (BZ#[1515352](#))

sysctl now allows tuned to reset kernel parameters

Prior to this update, a bug in **sysctl** did not allow kernel parameters to be set to default values. As a consequence, the **tuned** utility could not set default kernel parameters using **sysctl**. With this update, **sysctl** accepts default values to reset kernel parameters. As a result, **tuned** works as expected and kernel parameters can be reset to default values. (BZ#[1507356](#))

llvm-private no longer crashes when used together with more recent libstdc++ library versions

Previously, executable files in the **llvm-private** package providing drivers for graphics rendering were linked statically against the **libstdc++** library. As a consequence, running a program using GLX, the Mesa **llvmpipe** renderer, and a different **libstdc++** version caused an unexpected termination with message about **invalid pointer**. **llvm-private** has been changed and no longer statically links against **libstdc++**. As a result, programs using this driver no longer terminate unexpectedly in this situation. (BZ#[1417663](#))

ncat now correctly sets environment variables in UDP mode

Previously, the **ncat** utility did not set environment variables for User Datagram Protocol (UDP) connections properly. As a consequence, user's scripts failed in UDP mode. This update sets some internal values, and environment variables now are set properly. (BZ#[1573411](#))

ncat no longer uses the default HTTP port for all proxy types

Previously, the default port of an HTTP proxy was used even if another type of proxy, such as **socks4** or **socks5**, was specified. As a consequence, the **ncat** utility unsuccessfully tried to connect to a proxy type through the non-default port. This update corrects the code so that an HTTP proxy port is not used by default. As a result, **ncat** now sets the proper default port according to the proxy type. (BZ#[1546246](#))

Decoding and conversion of JPEG 2000 images now work correctly

Previously, decoding and conversion of JPEG 2000 images did not work correctly due to a bug in the **openjpeg** library. With this update, the underlying source code has been fixed, and decoding and conversion of JPEG 2000 images now works as expected. (BZ#[1553235](#))

strip no longer malforms binary files built with tools that use a later BFD library version

Previously, the **strip** tool created an invalid binary file if the file was originally produced by tools that use a later version of the **BFD** library than **strip**. As a consequence running the resulting binary file failed and generated an error message about an unresolvable relocation. **BFD** has been modified to report situations where it cannot recognize its future features instead of damaging the code that contains these features. As a result, **strip** now generates an error message and aborts in this situation. (BZ#[1553842](#))

CHAPTER 26. HARDWARE ENABLEMENT

The **lsslot -cpci** command now correctly reports PCI slot types

Prior to this update, the **lsslot -cpci** command reported PCI slot types as `<literal>Unknown slot type</literal>`. With this update, the bug has been fixed, and **lsslot** utility reports PCI slot types correctly. (BZ#1592429)

The **drmgr -c** command now loads the **rpادلpar_io** kernel module

Previously, the **drmgr -c** command did not select hotplug types. As a consequence, the **rpادلpar_io** kernel module was loaded only when one of the following hotplug types was explicitly selected: **pci**, **phb** or **slot**. With this update, the underlying source code has been fixed, and the command **drmgr -c** now loads **rpادلpar_io**. (BZ#1540067)

Diagnostic utilities now display CPU frequency values correctly

Due to a delay in the CPU cycle counter, diagnostic utilities, such as **lscpu**, in some cases reported an incorrect CPU frequency value. With this update, the affected utilities display correct values and report an error if the accurate value currently cannot be detected. (BZ#1596121)

The **ppc64_cpu** utility no longer fails when reading CPU frequency

Previously, the **ppc64_cpu** utility failed when reading CPU frequency. With this update, a thread is created for each of the CPUs if the number of CPUs in the system is less than **CPU_SETSIZE**, or only **CPU_SETSIZE** threads are created. As a result, the described problem no longer occurs, and **ppc64_cpu** no longer fails. (BZ#1628907)

CHAPTER 27. INSTALLATION AND BOOTING

The network service no longer hangs on stop or restart

Previously, when certain processes were executed from a network share, the **network** service could hang if it was stopped or restarted. A patch to the initscripts packages has been applied to not use the **pidof** utility, and the described problem no longer occurs. (BZ#1559384)

KSH no longer fails to process `/etc/init.d/functions`

The Korn Shell (KSH) is unable to process code where the word **local** appears on the same line as an array definition. This previously caused **KSH** to fail to source the `/etc/init.d/functions` file. This update provides a workaround to the **KSH** limitation, and the function file is now being sourced as expected.

Note that **KSH** may still be unable use some of the functions in `/etc/init.d/functions` file. This update only allows KSH to not fail during the sourcing of `/etc/init.d/functions`. (BZ#1554364)

Diskless NFS clients no longer hang when unmounting the root file system

Previously, diskless NFS clients became unresponsive in rare cases when the **network** service was stopped or restarted while unmounting the root file system. This happened because the unit files generated by **systemd** sometimes had incorrect dependencies.

A workaround has been applied in the initscripts package, and diskless NFS clients no longer hang in the described situation. (BZ#1572659)

A non-functioning `systemctl reload network.service` has been removed

The `systemctl reload network.service` command, which does not work due to technical limitations of initscripts has been removed, and using it now results in an appropriate warning message. To correctly apply a new configuration for the network service, use the **restart** command instead:

```
~]# systemctl restart network.service
```

(BZ#1554690)

Text mode will now prompt for a passphrase if a Kickstart file does not provide one while enabling encryption

Prior to this update, if you used the text mode interface with a Kickstart file that enabled disk encryption but did not provide a passphrase, the installation failed with an error. This update prompts the user to provide a passphrase during installation if the partitioning specified in the provided Kickstart file requires one. (BZ#1436304)

A `cmdline` Kickstart installation with conflicting packages now displays an error message

Previously, when a **cmdline** (noninteractive, unattended) Kickstart installation with conflicting packages was started, the installation failed and the machine rebooted before displaying the error message.

This update increases the reboot timeout from 10 to 180 seconds ensuring the appropriate error message is displayed. (BZ#1360223)

The custom partitioning screen now displays relevant storage configuration error messages

Previously, error messages in the custom partitioning screen were not always cleared after configuration changes. As a result, error messages that were not relevant to the current storage configuration were displayed.

This update ensures that the error messages displayed are relevant to the storage configuration in the custom partitioning screen. (BZ#1535781)

Host name is now configured correctly on an installed system

Previously, host name was not parsed properly from the IPv6 static configuration that was set by boot options. As a consequence, the host name specified by the **ip** installer boot option was not configured on an installed system.

The parsing of host name from the **ip** installer boot option has now been fixed for IPv6 static configuration. (BZ#1554271)

The **reqpart** Kickstart command will now only create partitions that are required by the hardware platform

Previously, when the **reqpart** command was specified in a Kickstart file and no partitions were required by the hardware platform, the installer attempted to perform automatic partitioning. As a result, the installation failed with an error.

This update ensures that the **reqpart** Kickstart command will only create partitions that are required by the hardware platform. (BZ#1557485)

Installation started with boot option **zfcpl.allow_lun_scan** is applied to the installed system

Previously, the boot option **zfcpl.allow_lun_scan** was not applied to the installed system and as a result, the installed system started without the boot option.

This update applies the boot option **zfcpl.allow_lun_scan** to the installed system. (BZ#1561662)

The **clearpart** Kickstart command can now be used on disk partitions

Previously, using the Kickstart command **clearpart --list=<part>** (where <part> is a partition on a disk) during installation worked for disks but not disk partitions.

As a consequence, Anaconda stopped the installation with the message:

```
Device <part> given in clearpart device list does not exist.
```

This update removes the restriction and supports clearing on disk partitions. (BZ#1561930)

CHAPTER 28. KERNEL

libcgroup no longer truncates the values of cgroup subsystem parameters that are longer than 100 characters

Previously, the internal representation of a value of any cgroup subsystem parameter was limited to have the length of 100 characters at maximum. Consequently, the **libcgroup** library truncated the values longer than 100 characters before writing them to a file representing matching cgroup subsystem parameter in the kernel. With this update, the maximal length of values of cgroup subsystem parameters in **libcgroup** has been extended to 4096 characters. As a result, **libcgroup** now handles values of cgroup subsystem parameters with any length correctly. (BZ#1549175)

The mlx5 device no longer contains a firmware issue

Previously, the **mlx5** device contained a firmware issue, which caused that the link of **mlx5** devices in certain situation dropped after rebooting a system. As a consequence, a message similar to the following was seen in the output of the **dmesg** command:

```
mlx5_core 0000:af:00.0: Port module event[error]: module 0, Cable error,  
Bus stuck(I2C or data shorted)
```

The issue is fixed in the latest firmware of this device. Contact your hardware vendor for information on how to obtain and install the latest firmware for your **mlx5** device. (BZ#1636930)

CHAPTER 29. REAL-TIME KERNEL

A race condition that prevented tasks from being scheduled properly has been fixed

Previously, preemption was enabled too early after a context switch. If a task was migrated to another CPU after a context switch, a mismatch between CPU and runqueue during load balancing sometimes occurred. Consequently, a runnable task on an idle CPU failed to run, and the operating system became unresponsive. This update disables preemption in the `schedule_tail()` function. As a result, CPU migration during post-schedule processing no longer occurs, which prevents the above mismatch. The operating system no longer hangs due to this bug. (BZ#[1608672](#), BZ#1541534)

CHAPTER 30. NETWORKING

Bad offload warnings are no longer displayed using `virtio_net`

Previously, using the `virtio_net` network adapter in bridge connections, user space programs sometimes generated Generic Segmentation Offload (GSO) packets with no checksum offload and passed them to the kernel. As a consequence, the kernel checksum offloading code displayed bad offload warnings unnecessarily. With this update, a patch has been applied, and the kernel does not warn anymore about bad checksum offload messages for such packets. (BZ#1544920)

The L2TP sequence number handling now works correctly

Previously, the kernel did not handle Layer 2 Tunneling Protocol (L2TP) sequence numbers properly and it was not compliant with RFC 3931. As a consequence, **L2TP** sessions stopped working unexpectedly. With this update, a patch has been applied to correctly handle sequence numbers in case of a packet loss. As a result, when users enable sequence numbers, **L2TP** sessions work as expected in the described scenario. (BZ#1527799)

The kernel no longer crashes when a `tunnel_key` mode is not specified

Previously, parsing configuration data in the `tunnel_key` action rules was incorrect if neither `set` nor `unset` mode was specified in the configuration. As a consequence, the kernel dereferenced an incorrect pointer and terminated unexpectedly. With this update, the kernel does not install `tunnel_key` if `set` or `unset` was not specified. As a result, the kernel no longer crashes in the described scenario. (BZ#1554907)

The `sysctl net.ipv4.route.min_pmtu` setting no longer set invalid values

Previously, the value provided by administrators for the `sysctl net.ipv4.route.min_pmtu` setting was not restricted. As a consequence, administrators were able to set a negative value for `net.ipv4.route.min_pmtu`. This sometimes resulted in setting the path Maximum Transmission Unit (MTU) of some routes to very large values because of an integer overflow. This update restricts values for `net.ipv4.route.min_pmtu` set to `>= 68`, the minimum valid MTU for IPv4. As a result, `net.ipv4.route.min_pmtu` can no longer be set to invalid values (negative value or `< 68`). (BZ#1541250)

`wpa_supplicant` no longer responds to packets whose destination address does not match the interface address

Previously, when `wpa_supplicant` was running on a Linux interface that was configured in `promiscuous` mode, incoming Extensible Authentication Protocol over LAN (EAPOL) packets were processed regardless of the destination address in the frame. However, `wpa_supplicant` checked the destination address only if the interface was enslaved to a bridge. As a consequence, in certain cases, `wpa_supplicant` was responding to EAPOL packets when the destination address was not the interface address. With this update, a socket filter has been added that allows the kernel to discard unicast EAPOL packets whose destination address does not match the interface address, and the described problem no longer occurs. (BZ#1434434)

`NetworkManager` no longer fails to detect duplicate IPv4 addresses

Previously, `NetworkManager` used to spawn an instance of the `arping` process to detect duplicate IPv4 addresses on the network. As a consequence, if the timeout configured for IPv4 Duplicate Address Detection (DAD) was short and the system was overloaded, `NetworkManager` sometimes failed to detect a duplicate address in time. With this update, the detection of duplicate IPv4 addresses is now performed internally to `NetworkManager` without spawning external binaries, and the described problem no longer occurs. (BZ#1507864)

`firewalld` now prevents partially applied rules

Previously, if a direct rule failed to be inserted for any reason, then all following direct rules with a higher

priority also failed to insert. As a consequence, direct rules were not applied completely. The processing has been changed to either apply all direct rules successfully or revert them all. As a result, if a rule failure occurs at startup, **firewalld** enters the **failed** status and allows the user to remedy the situation. This prevents unexpected results by having partially applied rules. (BZ#[1498923](#))

The wpa_supplicant upgrade no longer causes disconnections

Previously, the upgrade of the wpa_supplicant package caused a restart of the **wpa_supplicant** service. As a consequence, the network disconnected temporarily. With this update, the systemd unit is not restarted during the upgrade. As a result, the network connectivity no longer fails during the wpa_supplicant upgrade. (BZ#1505404)

CHAPTER 31. SECURITY

CardOS 5.3 smart cards with ECDSA support work correctly in OpenSC

Previously, OpenSC did not correctly parse the ECDSA algorithm in the **TokenInfo** information provided by CardOS 5.3 smart cards. As a consequence, OpenSC did not detect these cards. The **TokenInfo** parser has been updated and now complies with the PKCS #15 specification. As a result, CardOS 5.3 smart cards with ECDSA support work correctly in OpenSC. (BZ#[1562277](#))

Non-CCID-compliant smart card readers work in OpenSC

Certain smart card readers implement PIN pad functionality that does not follow the chip card interface device (CCID) specification. Previously, OpenSC detected the PIN pad of such smart card readers, but the reader could not be used with OpenSC. With this update, the PIN pad detection has been disabled in OpenSC by default. As a result, non-CCID-compliant smart card readers can be used, but without the PIN pad feature. (BZ#[1547117](#))

The **pkcs11-tool** utility now supports mechanism IDs and handles ECDSA keys correctly

Previously, the **pkcs11-tool** utility incorrectly handled **EC_POINT** values and support for certain vendor-specific mechanisms was missing. As a consequence, these mechanisms and certain ECDSA keys in hardware security modules (HSM) and smart cards were not supported by **pkcs11-tool**. With this update, the **pkcs11-tool** now handles **EC_POINT** values and vendor-specific mechanisms correctly. As a result, the utility now supports mechanism IDs and handles ECDSA keys correctly. (BZ#[1562572](#))

openscap RPM verification rules no longer work incorrectly with VM and container file systems

Previously, the **rpminfo**, **rpmverify**, and **rpmverifyfile** probes did not fully support offline mode. As a consequence, **OpenSCAP** RPM verification rules did not work correctly when scanning virtual machine (VM) and container file systems in offline mode. With this update, support for offline mode has been fixed, and results of scanning VM and container file systems in offline mode no longer contain false negatives. (BZ#[1556988](#))

sudo no longer blocks **poll()** for **/dev/ptmx**

Previously, when running a command through **sudo** that had the I/O logging enabled, a parent process of the command was occasionally blocked in the **poll()** function execution, waiting for an event on the **/dev/ptmx** file descriptor. Consequently, a deadlock occurred and **sudo** might leave the process of the command in an unresponsive state. This update adds a pseudoterminal cleanup logic, and **sudo** no longer causes a deadlock in the described scenario. (BZ#[1560657](#))

CHAPTER 32. SERVERS AND SERVICES

pxlcolor and pxlmono now work correctly

Previously, the **pxlcolor** and the **pxlmono** drivers in the Ghostscript interpreter did not function correctly. As a consequence, the drivers were likely to ignore a selection of a paper tray for certain printers, therefore only a specific paper tray was selected. This update applies a patch, which fixes the issue. As a result, the selection of different paper trays now works as expected in the described scenario. (BZ#1551782)

The nuxwdog service starts correctly when a sub-CA is installed

Previously, if a sub-CA was installed, the **nuxwdog** service did not allocate enough memory. As a consequence, the service failed to start. This update fixes the problem. As a result, **nuxwdog** starts correctly in the mentioned scenario. (BZ#[1615617](#))

Augeas reads /etc/fstab with white spaces more reliably

Previously, Augeas was not able to parse lines in the **/etc/fstab** file if they had white spaces at the beginning. This sometime caused problems in software tools that use Augeas, such as the **virt-v2v** utility or the **Puppet** management tool. With this update, the Fstab lens of Augeas correctly ignores white spaces at the beginning of lines. As a result, Augeas now reads **/etc/fstab** as expected. (BZ#[1544520](#))

CHAPTER 33. STORAGE

`mpathpersist` no longer fails when opening too many files

Previously, the `mpathpersist` utility sometimes overstepped the limit on open files when scanning a large number of devices. As a consequence, `mpathpersist` terminated unexpectedly.

With this update, `mpathpersist` now checks the `max_fds` configuration value and correctly sets the maximum number of open files. As a result, `mpathpersist` no longer fails when opening too many files. (BZ#1610263)

The `multipathd readsector0` checker now returns the correct result

Previously, in some cases the `multipathd` daemon was incorrectly calculating the I/O size to use with the `readsector0` checker, causing it to do a 0 size read. This could cause the `multipathd readsector0` checker to return the wrong result. It is also possible that some SCSI devices do not treat a 0 size read command as valid. With this fix, `multipathd` now uses the correct size for the `readsector0` checker. (BZ#1584228)

DM Multipath is much less likely to output an incorrect timeout error

Previously, Device Mapper Multipath (DM Multipath) printed an error message after reconfiguring devices for more than 10 seconds. The error was displayed even if the reconfigure was progressing successfully. As a consequence, it sometimes seemed that the reconfigure failed when reconfiguring a large number of devices.

With this update, the timeout limit has been increased from 10 to 60 seconds, and DM Multipath is now much less likely to print incorrect timeout errors when reconfiguring a large number of devices. (BZ#1544958)

`multipath` now correctly prints the `sysfs` state of paths

Previously, the `multipath -l` command did not print the `sysfs` state of paths because the `multipath` utility did not correctly set path information. With this update, the problem has been fixed, and `multipath` now prints the `sysfs` state of paths correctly. (BZ#1526876)

`multipathd` can now correctly set APTPL when registering keys on path devices

Previously, the `multipathd` service did not track which devices registered their persistent reservation keys with the Activate Persist Through Power Loss (APTPL) option. As a consequence, registrations always lost the APTPL setting.

With this update, the problem has been fixed:

- If you set the `reservation_key` option to a file in the `multipath.conf` configuration file, `multipathd` now keeps the APTPL setting automatically.
- If you set `reservation_key` to a specific key, you can now add the `:aptp1` string at the end of the key in `reservation_key`, which enables APTPL for it. Set this to match the APTPL setting used when registering the key. (BZ#1498724)

CHAPTER 34. SYSTEM AND SUBSCRIPTION MANAGEMENT

The `yum updateinfo` commands now respect `skip_if_unavailable` option

If a repository was configured with the `skip_if_unavailable=1` option, the `yum` commands operating on the `updateinfo` metadata, such as `yum updateinfo` or `yum check-update --security`, did not work correctly. Consequently, `yum` terminated with an error instead of skipping the repository. With this update, the underlying source code has been fixed to respect the `skip_if_unavailable` option. As a result, the affected `yum` commands now skip the unavailable repository as expected under the described circumstances. (BZ#1528608)

PART III. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 7.6.

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

CHAPTER 35. GENERAL UPDATES

The `systemd-importd` VM and container image import and export service

Latest `systemd` version now contains the `systemd-importd` daemon that was not enabled in the earlier build, which caused the `machinectl pull-*` commands to fail. Note that the `systemd-importd` daemon is offered as a Technology Preview and should not be considered stable.

(BZ#[1284974](#))

CHAPTER 36. AUTHENTICATION AND INTEROPERABILITY

Use of AD and LDAP sudo providers

The Active Directory (AD) provider is a back end used to connect to an AD server. Starting with Red Hat Enterprise Linux 7.2, using the AD sudo provider together with the LDAP provider is available as a Technology Preview. To enable the AD sudo provider, add the `sudo_provider=ad` setting in the `[domain]` section of the `sssd.conf` file. (BZ#1068725)

DNSSEC available as Technology Preview in IdM

Identity Management (IdM) servers with integrated DNS now support DNS Security Extensions (DNSSEC), a set of extensions to DNS that enhance security of the DNS protocol. DNS zones hosted on IdM servers can be automatically signed using DNSSEC. The cryptographic keys are automatically generated and rotated.

Users who decide to secure their DNS zones with DNSSEC are advised to read and follow these documents:

- DNSSEC Operational Practices, Version 2: <http://tools.ietf.org/html/rfc6781#section-2>
- Secure Domain Name System (DNS) Deployment Guide: <http://dx.doi.org/10.6028/NIST.SP.800-81-2>
- DNSSEC Key Rollover Timing Considerations: <http://tools.ietf.org/html/rfc7583>

Note that IdM servers with integrated DNS use DNSSEC to validate DNS answers obtained from other DNS servers. This might affect the availability of DNS zones that are not configured in accordance with recommended naming practices described in the Red Hat Enterprise Linux Networking Guide:

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Networking_Guide/ch-Configure_Host_Names.html#sec-Recommended_Naming_Practices. (BZ#1115294)

Identity Management JSON-RPC API available as Technology Preview

An API is available for Identity Management (IdM). To view the API, IdM also provides an API browser as Technology Preview.

In Red Hat Enterprise Linux 7.3, the IdM API was enhanced to enable multiple versions of API commands. Previously, enhancements could change the behavior of a command in an incompatible way. Users are now able to continue using existing tools and scripts even if the IdM API changes. This enables:

- Administrators to use previous or later versions of IdM on the server than on the managing client.
- Developers to use a specific version of an IdM call, even if the IdM version changes on the server.

In all cases, the communication with the server is possible, regardless if one side uses, for example, a newer version that introduces new options for a feature.

For details on using the API, see <https://access.redhat.com/articles/2728021> (BZ#1298286)

The Custodia secrets service provider is now available

As a Technology Preview, you can now use Custodia, a secrets service provider. Custodia stores or serves as a proxy for secrets, such as keys or passwords.

For details, see the upstream documentation at <http://custodia.readthedocs.io>. (BZ#1403214)

Containerized Identity Management server available as Technology Preview

The **rhel7/ipa-server** container image is available as a Technology Preview feature. Note that the **rhel7/sss** container image is now fully supported.

For details, see https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/using_containerized_identity_management_services. (BZ#1405325, BZ#1405326)

CHAPTER 37. CLUSTERING

The pcs tool now manages bundle resources in Pacemaker

As a Technology Preview starting with Red Hat Enterprise Linux 7.4, Pacemaker supports a special syntax for launching a Docker container with any infrastructure it requires: the bundle. After you have created a Pacemaker bundle, you can create a Pacemaker resource that the bundle encapsulates. For information on Pacemaker support for containers, see https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/high_availability_add-on_reference/.

There is one exception to this feature being Technology Preview: As of RHEL 7.4, Red Hat fully supports the usage of Pacemaker bundles for Red Hat Openstack Platform (RHOSP) deployments. (BZ#1433016)

New fence-agents-heuristics-ping fence agent

As a Technology Preview, Pacemaker now supports the **fence_heuristics_ping** agent. This agent aims to open a class of experimental fence agents that do no actual fencing by themselves but instead exploit the behavior of fencing levels in a new way.

If the heuristics agent is configured on the same fencing level as the fence agent that does the actual fencing but is configured before that agent in sequence, fencing issues an **off** action on the heuristics agent before it attempts to do so on the agent that does the fencing. If the heuristics agent gives a negative result for the **off** action it is already clear that the fencing level is not going to succeed, causing Pacemaker fencing to skip the step of issuing the **off** action on the agent that does the fencing. A heuristics agent can exploit this behavior to prevent the agent that does the actual fencing from fencing a node under certain conditions.

A user might want to use this agent, especially in a two-node cluster, when it would not make sense for a node to fence the peer if it can know beforehand that it would not be able to take over the services properly. For example, it might not make sense for a node to take over services if it has problems reaching the networking uplink, making the services unreachable to clients, a situation which a ping to a router might detect in that case. (BZ#1476401)

Heuristics supported in corosync-qdevice as a Technology Preview

Heuristics are a set of commands executed locally on startup, cluster membership change, successful connect to **corosync-qnetd**, and, optionally, on a periodic basis. When all commands finish successfully on time (their return error code is zero), heuristics have passed; otherwise, they have failed. The heuristics result is sent to **corosync-qnetd** where it is used in calculations to determine which partition should be quorate. (BZ#1413573, BZ#1389209)

New LVM and LVM lock manager resource agents

As a Technology Preview, Red Hat Enterprise Linux 7.6 introduces two new resource agents: **lvmlockd** and **LVM-activate**.

The **LVM-activate** agent provides a choice from multiple methods for LVM management throughout a cluster:

- tagging: the same as tagging with the existing **lvm** resource agent
- clvmd: the same as clvmd with the existing **lvm** resource agent
- system ID: a new option for using system ID for volume group failover (an alternative to tagging).
- lvmlockd: a new option for using **lvmlockd** and **dlm** for volume group sharing (an alternative to **clvmd**).

The new **lvmlockd** resource agent is used to start the **lvmlockd** daemon when **LVM-activate** is configured to use **lvmlockd**.

For information on the **lvmlockd** and **LVM-activate** resource agent, see the PCS help screens for those agents. For information on setting up LVM for use with **lvmlockd**, see the **lvmlockd(8)** man page. (BZ#[1513957](#), BZ#1634729)

CHAPTER 38. DESKTOP

Wayland available as a Technology Preview

The **Wayland** display server protocol is available in Red Hat Enterprise Linux as a Technology Preview with the dependent packages required to enable **Wayland** support in GNOME, which supports fractional scaling. **Wayland** uses the **libinput** library as its input driver.

The following features are currently unavailable or do not work correctly:

- Multiple GPU support is not possible at this time.
- The **NVIDIA** binary driver does not work under **Wayland**.
- The **xrandr** utility does not work under Wayland due to its different approach to handling, resolutions, rotations, and layout.
- Screen recording, remote desktop, and accessibility do not always work correctly under **Wayland**.
- No clipboard manager is available.
- It is currently impossible to restart **GNOME Shell** under **Wayland**.
- **Wayland** ignores keyboard grabs issued by X11 applications, such as virtual machines viewers. (BZ#1481411)

Fractional Scaling available as a Technology Preview

Starting with Red Hat Enterprise Linux 7.5, GNOME provides, as a Technology Preview, fractional scaling to address problems with monitors whose DPI lies in the middle between lo (scale 1) and hi (scale 2).

Due to technical limitations, fractional scaling is available only on Wayland. (BZ#[1481395](#))

CHAPTER 39. FILE SYSTEMS

ext4 and XFS file systems now support DAX

Starting with Red Hat Enterprise Linux 7.3, Direct Access (DAX) provides, as a Technology Preview, a means for an application to directly map persistent memory into its address space. To use DAX, a system must have some form of persistent memory available, usually in the form of one or more Non-Volatile Dual In-line Memory Modules (NVDIMMs), and a file system that supports DAX must be created on the NVDIMM(s). Also, the file system must be mounted with the **dax** mount option. Then, an **mmap** of a file on the dax-mounted file system results in a direct mapping of storage into the application's address space. (BZ#1274459)

pNFS block layout is now available

As a Technology Preview, Red Hat Enterprise Linux clients can now mount pNFS shares with the block layout feature.

Note that Red Hat recommends using the pNFS SCSI layout instead, which is similar to block layout but easier to use. (BZ#1111712)

OverlayFS

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

<https://www.kernel.org/doc/Documentation/filesystems/overlayfs.txt>.

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

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

- OverlayFS is only supported for use as a Docker graph driver. Its use can only be supported for container COW content, not for persistent storage. Any persistent storage must be placed on non-OverlayFS volumes to be supported. Only default Docker configuration can be used; that is, one level of overlay, one lowerdir, and both lower and upper levels are on the same file system.
- Only XFS is currently supported for use as a lower layer file system.
- On Red Hat Enterprise Linux 7.3 and earlier, SELinux must be enabled and in enforcing mode on the physical machine, but must be disabled in the container when performing container separation, that is the **/etc/sysconfig/docker** file must not contain **--selinux-enabled**. Starting with Red Hat Enterprise Linux 7.4, OverlayFS supports SELinux security labels, and you can enable SELinux support for containers by specifying **--selinux-enabled** in **/etc/sysconfig/docker**.
- The OverlayFS kernel ABI and userspace behavior are not considered stable, and may see changes in future updates.
- In order to make the yum and rpm utilities work properly inside the container, the user should be using the yum-plugin-ovl packages.

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

Note that XFS file systems must be created with the **-n ftype=1** option enabled for use as an overlay. With the rootfs and any file systems created during system installation, set the **--mkfsoptions=-n ftype=1** parameters in the Anaconda kickstart. When creating a new file system after the installation,

run the `# mkfs -t xfs -n ftype=1 /PATH/TO/DEVICE` command. To determine whether an existing file system is eligible for use as an overlay, run the `# xfs_info /PATH/TO/DEVICE | grep ftype` command to see if the `ftype=1` option is enabled.

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

<https://www.kernel.org/doc/Documentation/filesystems/overlayfs.txt>. (BZ#1206277)

Btrfs file system

The **Btrfs** (B-Tree) file system is available as a Technology Preview in Red Hat Enterprise Linux 7.

Red Hat Enterprise Linux 7.4 introduced the last planned update to this feature. **Btrfs** has been deprecated, which means Red Hat will not be moving **Btrfs** to a fully supported feature and it will be removed in a future major release of Red Hat Enterprise Linux. (BZ#1477977)

ima-evm-utils available as a Technology Preview for certain architectures

The `ima-evm-utils` package, available as a Technology Preview, provides utilities to label the file system and verify the integrity of your system at run time using the Integrity Measurement Architecture (IMA) and Extended Verification Module (EVM) features. These utilities enable you to monitor if files have been accidentally or maliciously altered.

Note that `ima-evm-utils` is now fully supported on the AMD64 and Intel 64 architecture, but remains in Technology Preview on all other architectures. (BZ#1384450)

CHAPTER 40. HARDWARE ENABLEMENT

LSI Syncro CS HA-DAS adapters

Red Hat Enterprise Linux 7.1 included code in the `megaraid_sas` driver to enable LSI Syncro CS high-availability direct-attached storage (HA-DAS) adapters. While the `megaraid_sas` driver is fully supported for previously enabled adapters, the use of this driver for Syncro CS is available as a Technology Preview. Support for this adapter is provided directly by LSI, your system integrator, or system vendor. Users deploying Syncro CS on Red Hat Enterprise Linux 7.2 and later are encouraged to provide feedback to Red Hat and LSI. For more information on LSI Syncro CS solutions, please visit <http://www.lsi.com/products/shared-das/pages/default.aspx>. (BZ#1062759)

tss2 enables TPM 2.0 for IBM Power LE

The `tss2` package adds IBM implementation of a Trusted Computing Group Software Stack (TSS) 2.0 as a Technology Preview for the IBM Power LE architecture. This package enables users to interact with TPM 2.0 devices. (BZ#1384452)

The `ibmvnic` device driver available as a Technology Preview

Since Red Hat Enterprise Linux 7.3, the IBM Virtual Network Interface Controller (vNIC) driver for IBM POWER architectures, `ibmvnic`, has been available as a Technology Preview. vNIC is a PowerVM virtual networking technology that delivers enterprise capabilities and simplifies network management. It is a high-performance, efficient technology that when combined with SR-IOV NIC provides bandwidth control Quality of Service (QoS) capabilities at the virtual NIC level. vNIC significantly reduces virtualization overhead, resulting in lower latencies and fewer server resources, including CPU and memory, required for network virtualization.

With Red Hat Enterprise Linux 7.6, the `ibmvnic` driver has been upgraded to version 1.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- The code that previously requested error information has been removed because no error ID is provided by the Virtual Input-Output (VIO) Server.
- Error reporting has been updated with the cause string. As a result, during a recovery, the driver classifies the string as a warning rather than an error.
- Error recovery on a login failure has been fixed.
- The failed state that occurred after a failover while migrating Logical Partitioning (LPAR) has been fixed.
- The driver can now handle all possible login response return values.
- A driver crash that happened during a failover or Link Power Management (LPM) if the Transmit and Receive (Tx/Rx) queues have changed has been fixed. (BZ#1519746)

CHAPTER 41. INSTALLATION AND BOOTING

Custom system image creation with Composer available as a Technology Preview

The Composer tool enables users to create customized RHEL images. Starting with Red Hat Enterprise Linux 7.6, Composer is available in the Extras channel as a Technology Preview in the `lorax-composer` package.

With Composer, users can create custom system images which include additional packages. Composer functionality can be accessed through a graphical user interface in Web Console, or with a command line interface in the **`composer-cli`** tool. Composer output formats include, among others:

- ISO disk image
- qcow2 file for direct use with a virtual machine
- file system image file

To install Composer, start it as a Web Console plugin and open the page in browser:

```
# yum install lorax-composer cockpit-composer composer-cli
# systemctl start lorax-composer
# systemctl start cockpit
$ gnome-open http://localhost:9090/
```

To automatically start Composer after reboot:

```
# systemctl enable lorax-composer.socket
# systemctl enable cockpit.socket
```

(BZ#1613966)

CHAPTER 42. KERNEL

Heterogeneous memory management included as a Technology Preview

Red Hat Enterprise Linux 7.3 introduced the heterogeneous memory management (HMM) feature as a Technology Preview. This feature has been added to the kernel as a helper layer for devices that want to mirror a process address space into their own memory management unit (MMU). Thus a non-CPU device processor is able to read system memory using the unified system address space. To enable this feature, add **experimental_hmm=enable** to the kernel command line. (BZ#1230959)

criu rebased to version 3.5

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

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

In Red Hat Enterprise Linux 7.6, the **criu** packages have been upgraded to upstream version 3.9, which provides a number of bug fixes and optimization for the runC container runtime. In addition, support for the 64-bit ARM architectures and the little-endian variant of IBM Power Systems CPU architectures has been fixed. (BZ#1400230, BZ#1464596)

kexec as a Technology Preview

The **kexec** system call has been provided as a Technology Preview. This system call enables loading and booting into another kernel from the currently running kernel, thus performing the function of the boot loader from within the kernel. Hardware initialization, which is normally done during a standard system boot, is not performed during a **kexec** boot, which significantly reduces the time required for a reboot. (BZ#1460849)

kexec fast reboot as a Technology Preview

The **kexec fast reboot** feature, which was introduced in Red Hat Enterprise Linux 7.5, continues to be available as a Technology Preview. **kexec fast reboot** makes the reboot significantly faster. To use this feature, you must load the **kexec** kernel manually, and then reboot the operating system. It is not possible to make **kexec fast reboot** as the default reboot action. Special case is using **kexec fast reboot** for **Anaconda**. It still does not enable to make **kexec fast reboot** default. However, when used with **Anaconda**, the operating system can automatically use **kexec fast reboot** after the installation is complete in case that user boots kernel with the **anaconda** option. To schedule a **kexec** reboot, use the **inst.kexec** command on the kernel command line, or include a **reboot --kexec** line in the Kickstart file. (BZ#1464377)

perf cqm has been replaced by resctrl

The Intel Cache Allocation Technology (CAT) was introduced in Red Hat Enterprise Linux 7.4 as a Technology Preview. However, the **perf cqm** tool did not work correctly due to an incompatibility between **perf** infrastructure and Cache Quality of Service Monitoring (CQM) hardware support. Consequently, multiple problems occurred when using **perf cqm**.

These problems included most notably:

- **perf cqm** did not support the group of tasks which is allocated using **resctrl**
- **perf cqm** gave random and inaccurate data due to several problems with recycling

- **perf cqm** did not provide enough support when running different kinds of events together (the different events are, for example, tasks, system-wide, and cgroup events)
- **perf cqm** provided only partial support for cgroup events
- The partial support for cgroup events did not work in cases with a hierarchy of cgroup events, or when monitoring a task in a cgroup and the cgroup together
- Monitoring tasks for the lifetime caused **perf** overhead
- **perf cqm** reported the aggregate cache occupancy or memory bandwidth over all sockets, while in most cloud and VMM-bases use cases the individual per-socket usage is needed

In Red Hat Enterprise Linux 7.5, **perf cqm** was replaced by the approach based on the **resctrl** file system, which addressed all of the aforementioned problems. (BZ#[1457533](#), BZ#1288964)

TC HW offloading available as a Technology Preview

Starting with Red Hat Enterprise Linux 7.6, Traffic Control (TC) Hardware offloading has been provided as a Technology Preview.

Hardware offloading enables that the selected functions of network traffic processing, such as shaping, scheduling, policing and dropping, are executed directly in the hardware instead of waiting for software processing, which improves the performance. (BZ#1503123)

AMD xgbe network driver available as a Technology Preview

Starting with Red Hat Enterprise Linux 7.6, the AMD **xgbe** network driver has been provided as a Technology Preview. (BZ#1589397)

CHAPTER 43. NETWORKING

Cisco usNIC driver

Cisco Unified Communication Manager (UCM) servers have an optional feature to provide a Cisco proprietary User Space Network Interface Controller (usNIC), which allows performing Remote Direct Memory Access (RDMA)-like operations for user-space applications. The `libusnic_verbs` driver, which is available as a Technology Preview, makes it possible to use usNIC devices via standard InfiniBand RDMA programming based on the Verbs API. (BZ#916384)

Cisco VIC kernel driver

The Cisco VIC Infiniband kernel driver, which is available as a Technology Preview, allows the use of Remote Directory Memory Access (RDMA)-like semantics on proprietary Cisco architectures. (BZ#916382)

Trusted Network Connect

Trusted Network Connect, available as a Technology Preview, is used with existing network access control (NAC) solutions, such as TLS, 802.1X, or IPsec to integrate endpoint posture assessment; that is, collecting an endpoint's system information (such as operating system configuration settings, installed packages, and others, termed as integrity measurements). Trusted Network Connect is used to verify these measurements against network access policies before allowing the endpoint to access the network. (BZ#755087)

SR-IOV functionality in the qlcn driver

Support for Single-Root I/O virtualization (SR-IOV) has been added to the `qlcn` driver as a Technology Preview. Support for this functionality will be provided directly by QLogic, and customers are encouraged to provide feedback to QLogic and Red Hat. Other functionality in the `qlcn` driver remains fully supported. (BZ#1259547)

The `flower` classifier with off-loading support

`flower` is a Traffic Control (TC) classifier intended to allow users to configure matching on well-known packet fields for various protocols. It is intended to make it easier to configure rules over the `u32` classifier for complex filtering and classification tasks. `flower` also supports the ability to off-load classification and action rules to underlying hardware if the hardware supports it. The `flower` TC classifier is now provided as a Technology Preview. (BZ#1393375)

CHAPTER 44. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE

The `postfix` role of Red Hat Enterprise Linux System Roles as a Technology Preview

Red Hat Enterprise Linux System Roles provides a configuration interface for Red Hat Enterprise Linux subsystems, which makes system configuration easier through the inclusion of Ansible Roles. This interface enables managing system configurations across multiple versions of Red Hat Enterprise Linux, as well as adopting new major releases.

Since Red Hat Enterprise Linux 7.4, the Red Hat Enterprise Linux System Roles packages have been distributed through the Extras channel. For details regarding Red Hat Enterprise Linux System Roles, see <https://access.redhat.com/articles/3050101>.

Red Hat Enterprise Linux System Roles currently consists of five roles:

- `selinux`
- `kdump`
- `network`
- `timesync`
- `postfix`

The `postfix` role has been available as a Technology Preview since Red Hat Enterprise Linux 7.4.

The remaining roles have been fully supported since Red Hat Enterprise Linux 7.6. (BZ#1439896)

CHAPTER 45. SECURITY

USBGuard enables blocking USB devices while the screen is locked as a Technology Preview

With the **USBGuard** framework, you can influence how an already running **usbguard-daemon** instance handles newly inserted USB devices by setting the value of the **InsertedDevicePolicy** runtime parameter. This functionality is provided as a Technology Preview, and the default choice is to apply the policy rules to figure out whether to authorize the device or not.

See the **Blocking USB devices while the screen is locked** Knowledge Base article: <https://access.redhat.com/articles/3230621> (BZ#1480100)

pk12util can now import certificates signed with RSA-PSS

The **pk12util** tool now provides importing a certificate signed with the **RSA-PSS** algorithm as a Technology Preview.

Note that if the corresponding private key is imported and has the **PrivateKeyInfo.privateKeyAlgorithm** field that restricts the signing algorithm to **RSA-PSS**, it is ignored when importing the key to a browser. See https://bugzilla.mozilla.org/show_bug.cgi?id=1413596 for more information. (BZ#1431210)

Support for certificates signed with RSA-PSS in certutil has been improved

Support for certificates signed with the **RSA-PSS** algorithm in the **certutil** tool has been improved. Notable enhancements and fixes include:

- The **--pss** option is now documented.
- The **PKCS#1 v1.5** algorithm is no longer used for self-signed signatures when a certificate is restricted to use **RSA-PSS**.
- Empty **RSA-PSS** parameters in the **subjectPublicKeyInfo** field are no longer printed as invalid when listing certificates.
- The **--pss-sign** option for creating regular RSA certificates signed with the **RSA-PSS** algorithm has been added.

Support for certificates signed with **RSA-PSS** in **certutil** is provided as a Technology Preview. (BZ#1425514)

NSS is now able to verify RSA-PSS signatures on certificates

With the new version of the **nss** package, the **Network Security Services** (NSS) libraries now provide verifying **RSA-PSS** signatures on certificates as a Technology Preview. Prior to this update, clients using **NSS** as the **SSL** backend were not able to establish a **TLS** connection to a server that offered only certificates signed with the **RSA-PSS** algorithm.

Note that the functionality has the following limitations:

- The algorithm policy settings in the **/etc/pki/nss-legacy/rhel7.config** file do not apply to the hash algorithms used in **RSA-PSS** signatures.
- **RSA-PSS** parameters restrictions between certificate chains are ignored and only a single certificate is taken into account. (BZ#1432142)

SECCOMP can be now enabled in libreswan

As a Technology Preview, the **seccomp=enabled|tolerant|disabled** option has been added to the **ipsec.conf** configuration file, which makes it possible to use the Secure Computing mode (SECCOMP). This improves the syscall security by whitelisting all the system calls that **Libreswan** is allowed to execute. For more information, see the **ipsec.conf(5)** man page. (BZ#[1375750](#))

CHAPTER 46. STORAGE

Multi-queue I/O scheduling for SCSI

Red Hat Enterprise Linux 7 includes a new multiple-queue I/O scheduling mechanism for block devices known as blk-mq. The `scsi-mq` package allows the Small Computer System Interface (SCSI) subsystem to make use of this new queuing mechanism. This functionality is provided as a Technology Preview and is not enabled by default. To enable it, add `scsi_mod.use_blk_mq=Y` to the kernel command line.

Although blk-mq is intended to offer improved performance, particularly for low-latency devices, it is not guaranteed to always provide better performance. In particular, in some cases, enabling `scsi-mq` can result in significantly worse performance, especially on systems with many CPUs. (BZ#1109348)

Targetd plug-in from the libStorageMgmt API

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

The Targetd plug-in is not fully supported and remains a Technology Preview. (BZ#1119909)

Support for Data Integrity Field/Data Integrity Extension (DIF/DIX)

DIF/DIX is a new addition to the SCSI Standard. It is fully supported in Red Hat Enterprise Linux 7 for the HBAs and storage arrays specified in the Features chapter, but it remains in Technology Preview for all other HBAs and storage arrays.

DIF/DIX increases the size of the commonly used 512 byte disk block from 512 to 520 bytes, adding the Data Integrity Field (DIF). The DIF stores a checksum value for the data block that is calculated by the Host Bus Adapter (HBA) when a write occurs. The storage device then confirms the checksum on receipt, and stores both the data and the checksum. Conversely, when a read occurs, the checksum can be verified by the storage device, and by the receiving HBA. (BZ#1072107)

SCSI-MQ as a Technology Preview in the `q1a2xxx` and `lpfc` drivers

The `q1a2xxx` driver updated in Red Hat Enterprise Linux 7.4 can enable the use of SCSI-MQ (multiqueue) with the `q12xmqsupport=1` module parameter. The default value is `0` (disabled).

The SCSI-MQ functionality is provided as a Technology Preview when used with the `q1a2xxx` or the `lpfc` drivers.

Note that a recent performance testing at Red Hat with async IO over Fibre Channel adapters using SCSI-MQ has shown significant performance degradation under certain conditions. (BZ#1414957)

NVMe/FC available as a Technology Preview in Qlogic adapters using the `q1a2xxx` driver

The NVMe over Fibre Channel (NVMe/FC) transport type is available as a Technology Preview in Qlogic adapters using the `q1a2xxx` driver.

NVMe/FC is an additional fabric transport type for the Nonvolatile Memory Express (NVMe) protocol, in addition to the Remote Direct Memory Access (RDMA) protocol that was previously introduced in Red Hat Enterprise Linux.

NVMe/FC provides a higher-performance, lower-latency I/O protocol over existing Fibre Channel infrastructure. This is especially important with solid-state storage arrays, because it allows the performance benefits of NVMe storage to be passed through the fabric transport, rather than being encapsulated in a different protocol, SCSI.

Note that since Red Hat Enterprise Linux 7.6, NVMe/FC is fully supported with Broadcom Emulex Fibre Channel 32Gbit adapters using the **lpfc** driver. See the restrictions listed in the New Features part. (BZ#[1387768](#), BZ#1454386)

CHAPTER 47. SYSTEM AND SUBSCRIPTION MANAGEMENT

YUM 4 available as Technology Preview

YUM version 4, a next generation of the YUM package manager, is now available as a Technology Preview in the Red Hat Enterprise Linux 7 Extras channel.

YUM 4 is based on the **DNF** technology and offers the following advantages over the standard **YUM 3** used on RHEL 7:

- Increased performance
- Support for modular content
- Well-designed stable API for integration with tooling

To install **YUM 4**, run the **yum install nextgen-yum4** command.

Make sure to install the **dnf-plugin-subscription-manager** package, which includes the **subscription-manager** plug-in. This plug-in is required for accessing protected repositories provided by the Red Hat Customer Portal or Red Hat Satellite 6, and for automatic updates of the **/etc/yum.repos.d/redhat.repo** file.

To manage packages, use the **yum4** command and its particular options the same way as the **yum** command.

For detailed information about differences between the new **YUM 4** tool and **YUM 3**, see http://dnf.readthedocs.io/en/latest/cli_vs_yum.html. (BZ#1461652, BZ#1558411)

CHAPTER 48. VIRTUALIZATION

eBPF system call for tracing

Red Hat Enterprise Linux 7.6 introduces the Extended Berkeley Packet Filter tool (eBPF) as a Technology Preview. This tool is enabled only for the tracing subsystem. For details, see the Red Hat Knowledgebase article at <https://access.redhat.com/articles/3550581>. (BZ#1559615, BZ#1559756, BZ#1311586)

USB 3.0 support for KVM guests

USB 3.0 host adapter (xHCI) emulation for KVM guests remains a Technology Preview in Red Hat Enterprise Linux 7. (BZ#1103193)

Select Intel network adapters now support SR-IOV as a guest on Hyper-V

In this update for Red Hat Enterprise Linux guest virtual machines running on Hyper-V, a new PCI passthrough driver adds the ability to use the single-root I/O virtualization (SR-IOV) feature for Intel network adapters supported by the ixgbev driver. This ability is enabled when the following conditions are met:

- SR-IOV support is enabled for the network interface controller (NIC)
- SR-IOV support is enabled for the virtual NIC
- SR-IOV support is enabled for the virtual switch

The virtual function (VF) from the NIC is attached to the virtual machine.

The feature is currently supported with Microsoft Windows Server 2016. (BZ#1348508)

No-IOMMU mode for VFIO drivers

As a Technology Preview, this update adds No-IOMMU mode for virtual function I/O (VFIO) drivers. The No-IOMMU mode provides the user with full user-space I/O (UIO) access to a direct memory access (DMA)-capable device without a I/O memory management unit (IOMMU). Note that in addition to not being supported, using this mode is not secure due to the lack of I/O management provided by IOMMU. (BZ#1299662)

virt-v2v can now use vmx configuration files to convert VMware guests

As a Technology Preview, the **virt-v2v** utility now includes the **vmx** input mode, which enables the user to convert a guest virtual machine from a VMware vmx configuration file. Note that to do this, you also need access to the corresponding VMware storage, for example by mounting the storage using NFS. It is also possible to access the storage using SSH, by adding the **-it ssh** parameter. (BZ#1441197, BZ#1523767)

virt-v2v can convert Debian and Ubuntu guests

As a technology preview, the **virt-v2v** utility can now convert Debian and Ubuntu guest virtual machines. Note that the following problems currently occur when performing this conversion:

- **virt-v2v** cannot change the default kernel in the GRUB2 configuration, and the kernel configured in the guest is not changed during the conversion, even if a more optimal version of the kernel is available on the guest.
- After converting a Debian or Ubuntu VMware guest to KVM, the name of the guest's network interface may change, and thus requires manual configuration. (BZ#1387213)

Virtio devices can now use vIOMMU

As a Technology Preview, this update enables virtio devices to use virtual Input/Output Memory

Management Unit (viOMMU). This guarantees the security of Direct Memory Access (DMA) by allowing the device to DMA only to permitted addresses. However, note that only guest virtual machines using Red Hat Enterprise Linux 7.4 or later are able to use this feature. (BZ#[1283251](#), BZ#1464891)

virt-v2v converts VMWare guests faster and more reliably

As a Technology Preview, the **virt-v2v** utility can now use the VMWare Virtual Disk Development Kit (VDDK) to import a VMWare guest virtual machine to a KVM guest. This enables **virt-v2v** to connect directly to the VMWare ESXi hypervisor, which improves the speed and reliability of the conversion.

Note that this conversion import method requires the external **nbdkit** utility and its VDDK plug-in. (BZ#1477912)

Open Virtual Machine Firmware

The Open Virtual Machine Firmware (OVMF) is available as a Technology Preview in Red Hat Enterprise Linux 7. OVMF is a UEFI secure boot environment for AMD64 and Intel 64 guests. (BZ#653382)

GPU-based mediated devices now support the VNC console

As a Technology Preview, the Virtual Network Computing (VNC) console is now available for use with GPU-based mediated devices, such as the NVIDIA vGPU technology. As a result, it is now possible to use these mediated devices for real-time rendering of a virtual machine's graphical output. (BZ#[1475770](#), BZ#1470154, BZ#1555246)

PART IV. DEVICE DRIVERS

This part provides a comprehensive listing of all device drivers that are new or have been updated in Red Hat Enterprise Linux 7.6.

CHAPTER 49. NEW DRIVERS

Network Drivers

- Thunderbolt network driver (thunderbolt-net.ko.xz).
- AMD 10 Gigabit Ethernet Driver (amd-xgbe.ko.xz).

Storage Drivers

- Command Queue Host Controller Interface driver (cqhci.ko.xz).

Graphics Drivers and Miscellaneous Drivers

- DRM GPU scheduler (gpu-sched.ko.xz).
- Closed hash table (chash.ko.xz).
- RMI4 SMBus driver (rmi_smbus.ko.xz).
- RMI bus.
- RMI F03 module (rmi_core.ko.xz).
- Dell WMI descriptor driver (dell-wmi-descriptor.ko.xz).
- Intel® PMC Core Driver (intel_pmc_core.ko.xz).
- Intel® WMI Thunderbolt force power driver (intel-wmi-thunderbolt.ko.xz).
- ACPI Hardware Watchdog (WDAT) driver (wdat_wdt.ko.xz).
- IIO helper functions for setting up triggered buffers (industrialio-triggered-buffer.ko.xz).
- HID Sensor Pressure (hid-sensor-pressure.ko.xz).
- HID Sensor Device Rotation (hid-sensor-rotation.ko.xz).
- HID Sensor Inclinator 3D (hid-sensor-incl-3d.ko.xz).
- HID Sensor trigger processing (hid-sensor-trigger.ko.xz).
- HID Sensor common attribute processing (hid-sensor-iio-common.ko.xz).
- HID Sensor Magnetometer 3D (hid-sensor-magn-3d.ko.xz).
- HID Sensor ALS (hid-sensor-als.ko.xz).
- HID Sensor Proximity (hid-sensor-prox.ko.xz).
- HID Sensor Gyroscope 3D (hid-sensor-gyro-3d.ko.xz).
- HID Sensor Accel 3D (hid-sensor-accel-3d.ko.xz).
- HID Sensor Hub driver (hid-sensor-hub.ko.xz).
- HID Sensor Custom and Generic sensor driver (hid-sensor-custom.ko.xz).

CHAPTER 50. UPDATED DRIVERS

Storage Driver Updates

- The Microsemi Smart Family Controller driver (smartpqi.ko.xz) has been updated to version 1.1.4-115.
- The HP Smart Array Controller driver (hpsa.ko.xz) has been updated to version 3.4.20-125-RH1.
- The Emulex LightPulse Fibre Channel SCSI driver (lpfc.ko.xz) has been updated to version 0:12.0.0.5.
- The Avago MegaRAID SAS driver (megaraid_sas.ko.xz) has been updated to version 07.705.02.00-rh1.
- The Dell PERC2, 2/Si, 3/Si, 3/Di, Adaptec Advanced Raid Products, HP NetRAID-4M, IBM ServeRAID & ICP SCS driver (aacraid.ko.xz) has been updated to version 1.2.1[50877]-custom.
- The QLogic FastLinQ 4xxxx iSCSI Module driver (qedf.ko.xz) has been updated to version 8.33.0.20.
- The QLogic Fibre Channel HBA driver (qla2xxx.ko.xz) has been updated to version 10.00.00.06.07.6-k.
- The QLogic QEDF 25/40/50/100Gb FCoE driver (qedf.ko.x) has been updated to version 8.33.0.20.
- The LSI MPT Fusion SAS 3.0 Device driver (mpt3sas.ko.xz) has been updated to version 16.100.01.00.
- The LSI MPT Fusion SAS 2.0 Device driver (mpt2sas.ko.xz) has been updated to version 20.103.01.00.

Network Driver Updates

- The Realtek RTL8152/RTL8153 Based USB Ethernet Adapters driver (r8152.ko.xz) has been updated to version v1.09.9.
- The VMware vmxnet3 virtual NIC driver (vmxnet3.ko.xz) has been updated to version 1.4.14.0-k.
- The Intel® Ethernet Connection XL710 Network driver (i40e.ko.xz) has been updated to version 2.3.2-k.
- The Intel® 10 Gigabit Virtual Function Network driver (ixgbevf.ko.xz) has been updated to version 4.1.0-k-rh7.6.
- The Intel® 10 Gigabit PCI Express Network driver (ixgbe.ko.xz) has been updated to version 5.1.0-k-rh7.6.
- The Intel® XL710 X710 Virtual Function Network driver (i40evf.ko.xz) has been updated to version 3.2.2-k.
- The Intel® Ethernet Switch Host Interface driver (fm10k.ko.xz) has been updated to version 0.22.1-k.
- The Broadcom BCM573xx network driver (bnxt_en.ko.xz) has been updated to version 1.9.1.

- The Cavium LiquidIO Intelligent Server Adapter driver (liquidio.ko.xz) has been updated to version 1.7.2.
- The Cavium LiquidIO Intelligent Server Adapter Virtual Function driver (liquidio_vf.ko.xz) has been updated to version 1.7.2.
- The Elastic Network Adapter (ENA) driver (ena.ko.xz) has been updated to version 1.5.0K.
- The aQuantia Corporation Network driver (atlantic.ko.xz) has been updated to version 2.0.2.1-kern.
- The QLogic FastLinQ 4xxxx Ethernet driver (qed.ko.xz) has been updated to version 8.33.0.20.
- The QLogic FastLinQ 4xxxx Core Module driver (qed.ko.xz) has been updated to version 8.33.0.20.
- The Cisco VIC Ethernet NIC driver (enic.ko.xz) has been updated to version 2.3.0.53.

Graphics Driver and Miscellaneous Driver Updates

- The VMware Memory Control (Balloon) driver (vmw_balloon.ko.xz) has been updated to version 1.4.1.0-k.
- The HP watchdog driver (hpwdt.ko.xz) has been updated to version 1.4.0-RH1k.
- The standalone drm driver for the VMware SVGA device (vmwgfx.ko.xz) has been updated to version 2.14.1.0.

CHAPTER 51. DEPRECATED FUNCTIONALITY

This chapter provides an overview of functionality that has been deprecated in all minor releases of Red Hat Enterprise Linux 7 up to Red Hat Enterprise Linux 7.6.

Deprecated functionality continues to be supported until the end of life of Red Hat Enterprise Linux 7. 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.

Python 2 has been deprecated

Python 2 will be replaced with **Python 3** in the next Red Hat Enterprise Linux (RHEL) major release.

See the [Conservative Python 3 Porting Guide](#) for information on how to migrate large code bases to **Python 3**.

Note that **Python 3** is available to RHEL customers, and supported on RHEL, as a part of [Red Hat Software Collections](#).

LVM libraries and LVM Python bindings have been deprecated

The **lvm2app** library and LVM Python bindings, which are provided by the `lvm2-python-libs` package, have been deprecated.

Red Hat recommends the following solutions instead:

- The LVM D-Bus API in combination with the **lvm2-dbusd** service. This requires using Python version 3.
- The LVM command-line utilities with JSON formatting; this formatting has been available since the `lvm2` package version 2.02.158.

Mirrored mirror log has been deprecated in LVM

The mirrored mirror log feature of mirrored LVM volumes has been deprecated. A future major release of Red Hat Enterprise Linux will no longer support creating or activating LVM volumes with a mirrored mirror log.

The recommended replacements are:

- RAID1 LVM volumes. The main advantage of RAID1 volumes is their ability to work even in degraded mode and to recover after a transient failure. For information on converting mirrored volumes to RAID1, see the [Converting a Mirrored LVM Device to a RAID1 Device](#) section in the LVM Administration guide.
- Disk mirror log. To convert a mirrored mirror log to disk mirror log, use the following command:
lvconvert --mirrorlog disk my_vg/my_lv.

Deprecated packages related to Identity Management and security

The following packages have been deprecated and will not be included in a future major release of Red Hat Enterprise Linux:

Deprecated packages	Proposed replacement package or product
authconfig	authselect
pam_pkcs11	sssd [a]
pam_krb5	sssd [b]
openldap-servers	Depending on the use case, migrate to Identity Management included in Red Hat Enterprise Linux or to Red Hat Directory Server. [c]
mod_auth_kerb	mod_auth_gssapi
python-kerberos python-krbV	python-gssapi
python-requests-kerberos	python-requests-gssapi
hesiod	No replacement available.
mod_nss	mod_ssl
mod_revocator	No replacement available.
<p>[a] System Security Services Daemon (SSSD) contains enhanced smart card functionality.</p> <p>[b] For details on migrating from pam_krb5 to sssd, see Migrating from pam_krb5 to sssd in the upstream SSSD documentation.</p> <p>[c] Red Hat Directory Server requires a valid Directory Server subscription. For details, see also What is the support status of the LDAP-server shipped with Red Hat Enterprise Linux? in Red Hat Knowledgebase.</p>	



NOTE

In Red Hat Enterprise Linux 7.5, the following packages were added to the table above:

- `mod_auth_kerb`
- `python-kerberos`, `python-krbV`
- `python-requests-kerberos`
- `hesiod`
- `mod_nss`
- `mod_revocator`

The Clevis HTTP pin has been deprecated

The **Clevis** HTTP pin has been deprecated and this feature will not be included in the next major version of Red Hat Enterprise Linux and will remain out of the distribution until a further notice.

3DES is removed from the Python SSL default cipher list

The Triple Data Encryption Standard (**3DES**) algorithm has been removed from the **Python** SSL default cipher list. This enables **Python** applications using SSL to be PCI DSS-compliant.

sssd-secrets has been deprecated

The **sssd-secrets** component of the **System Security Services Daemon** (SSSD) has been deprecated in Red Hat Enterprise Linux 7.6. This is because Custodia, a secrets service provider, is no longer actively developed. Use other Identity Management tools to store secrets, for example the Vaults.

Support for earlier IdM servers and for IdM replicas at domain level 0 will be limited

Red Hat does not plan to support using Identity Management (IdM) servers running Red Hat Enterprise Linux (RHEL) 7.3 and earlier with IdM clients of the next major release of RHEL. If you plan to introduce client systems running on the next major version of RHEL into a deployment that is currently managed by IdM servers running on RHEL 7.3 or earlier, be aware that you will need to upgrade the servers, moving them to RHEL 7.4 or later.

In the next major release of RHEL, only domain level 1 replicas will be supported. Before introducing IdM replicas running on the next major version of RHEL into an existing deployment, be aware that you will need to upgrade all IdM servers to RHEL 7.4 or later, and change the domain level to 1.

Consider planning the upgrade in advance if your deployment will be affected.

Bug-fix only support for the nss-pam-ldapd and NIS packages in the next major release of Red Hat Enterprise Linux

The `nss-pam-ldapd` packages and packages related to the **NIS server** will be released in the future major release of Red Hat Enterprise Linux but will receive a limited scope of support. Red Hat will accept bug reports but no new requests for enhancements. Customers are advised to migrate to the following replacement solutions:

Affected packages	Proposed replacement package or product
<code>nss-pam-ldapd</code>	<code>sssd</code>

Affected packages	Proposed replacement package or product
ypserv ypbind portmap yp-tools	Identity Management in Red Hat Enterprise Linux

Use the Go Toolset instead of golang

The golang package, previously available in the Optional channel, will no longer receive updates in Red Hat Enterprise Linux 7. Developers are encouraged to use the **Go Toolset** instead, which is available through the [Red Hat Developer program](#).

mesa-private-llvm will be replaced with llvm-private

The mesa-private-llvm package, which contains the LLVM-based runtime support for **Mesa**, will be replaced in a future minor release of Red Hat Enterprise Linux 7 with the llvm-private package.

libdbi and libdbi-drivers have been deprecated

The libdbi and libdbi-drivers packages will not be included in the next Red Hat Enterprise Linux (RHEL) major release.

Ansible deprecated in the Extras channel

Ansible and its dependencies will no longer be updated through the Extras channel. Instead, the Red Hat Ansible Engine product has been made available to Red Hat Enterprise Linux subscriptions and will provide access to the official Ansible Engine channel. Customers who have previously installed **Ansible** and its dependencies from the Extras channel are advised to enable and update from the Ansible Engine channel, or uninstall the packages as future errata will not be provided from the Extras channel.

Ansible was previously provided in Extras (for AMD64 and Intel 64 architectures, and IBM POWER, little endian) as a runtime dependency of, and limited in support to, the Red Hat Enterprise Linux (RHEL) System Roles. Ansible Engine is available today for AMD64 and Intel 64 architectures, with IBM POWER, little endian availability coming soon.

Note that **Ansible** in the Extras channel was not a part of the Red Hat Enterprise Linux FIPS validation process.

The following packages have been deprecated from the Extras channel:

- ansible(-doc)
- libtomcrypt
- libtommath(-devel)
- python2-crypto
- python2-jmespath
- python-httplib2
- python-paramiko(-doc)
- python-passlib

- `sshpas`

For more information and guidance, see the Knowledgebase article at <https://access.redhat.com/articles/3359651>.

Note that Red Hat Enterprise Linux System Roles continue to be distributed through the Extras channel. Although Red Hat Enterprise Linux System Roles no longer depend on the `ansible` package, installing `ansible` from the Ansible Engine repository is still needed to run playbooks which use Red Hat Enterprise Linux System Roles.

`signtool` has been deprecated

The **`signtool`** tool from the `nss` packages, which uses insecure signature algorithms, has been deprecated and will not be included in a future minor release of Red Hat Enterprise Linux.

TLS compression support has been removed from `nss`

To prevent security risks, such as the CRIME attack, support for TLS compression in the **`NSS`** library has been removed for all TLS versions. This change preserves the API compatibility.

Public web CAs are no longer trusted for code signing by default

The Mozilla CA certificate trust list distributed with Red Hat Enterprise Linux 7.5 no longer trusts any public web CAs for code signing. As a consequence, any software that uses the related flags, such as **`NSS`** or **`OpenSSL`**, no longer trusts these CAs for code signing by default. The software continues to fully support code signing trust. Additionally, it is still possible to configure CA certificates as trusted for code signing using system configuration.

`Sendmail` has been deprecated

`Sendmail` has been deprecated in Red Hat Enterprise Linux 7. Customers are advised to use **`Postfix`**, which is configured as the default Mail Transfer Agent (MTA).

`dmraid` has been deprecated

Since Red Hat Enterprise Linux 7.5, the `dmraid` packages have been deprecated. It will stay available in Red Hat Enterprise Linux 7 releases but a future major release will no longer support legacy hybrid combined hardware and software RAID host bus adapter (HBA).

Automatic loading of `DCCP` modules through socket layer is now disabled by default

For security reasons, automatic loading of the **`Datagram Congestion Control Protocol (DCCP)`** kernel modules through socket layer is now disabled by default. This ensures that userspace applications can not maliciously load any modules. All **`DCCP`** related modules can still be loaded manually through the **`modprobe`** program.

The `/etc/modprobe.d/dccp-blacklist.conf` configuration file for blacklisting the **`DCCP`** modules is included in the kernel package. Entries included there can be cleared by editing or removing this file to restore the previous behavior.

Note that any re-installation of the same kernel package or of a different version does not override manual changes. If the file is manually edited or removed, these changes persist across package installations.

`rsyslog-libdbi` has been deprecated

The `rsyslog-libdbi` sub-package, which contains one of the less used **`rsyslog`** module, has been deprecated and will not be included in a future major release of Red Hat Enterprise Linux. Removing unused or rarely used modules helps users to conveniently find a database output to use.

The `inputname` option of the `rsyslog imudp` module has been deprecated

The **inputname** option of the **imudp** module for the **rsyslog** service has been deprecated. Use the **name** option instead.

SMBv1 is no longer installed with Microsoft Windows 10 and 2016 (updates 1709 and later)

Microsoft announced that the Server Message Block version 1 (SMBv1) protocol will no longer be installed with the latest versions of Microsoft Windows and Microsoft Windows Server. Microsoft also recommends users to disable SMBv1 on earlier versions of these products.

This update impacts Red Hat customers who operate their systems in a mixed Linux and Windows environment. Red Hat Enterprise Linux 7.1 and earlier support only the SMBv1 version of the protocol. Support for SMBv2 was introduced in Red Hat Enterprise Linux 7.2.

For details on how this change affects Red Hat customers, see [SMBv1 no longer installed with latest Microsoft Windows 10 and 2016 update \(version 1709\)](#) in Red Hat Knowledgebase.

FedFS has been deprecated

Federated File System (FedFS) has been deprecated because the upstream FedFS project is no longer being actively maintained. Red Hat recommends migrating FedFS installations to use **autofs**, which provides more flexible functionality.

Btrfs has been deprecated

The **Btrfs** file system has been in Technology Preview state since the initial release of Red Hat Enterprise Linux 6. Red Hat will not be moving **Btrfs** to a fully supported feature and it will be removed in a future major release of Red Hat Enterprise Linux.

The **Btrfs** file system did receive numerous updates from the upstream in Red Hat Enterprise Linux 7.4 and will remain available in the Red Hat Enterprise Linux 7 series. However, this is the last planned update to this feature.

tcp_wrappers deprecated

The **tcp_wrappers** package has been deprecated. **tcp_wrappers** provides a library and a small daemon program that can monitor and filter incoming requests for audit, cyrus-imap, dovecot, nfs-utils, openssh, openldap, proftpd, sendmail, stunnel, syslog-ng, vsftpd, and various other network services.

nautilus-open-terminal replaced with gnome-terminal-nautilus

Since Red Hat Enterprise Linux 7.3, the **nautilus-open-terminal** package has been deprecated and replaced with the **gnome-terminal-nautilus** package. This package provides a Nautilus extension that adds the **Open in Terminal** option to the right-click context menu in Nautilus. **nautilus-open-terminal** is replaced by **gnome-terminal-nautilus** during the system upgrade.

sslwrap() removed from Python

The **sslwrap()** function has been removed from **Python 2.7**. After the [466 Python Enhancement Proposal](#) was implemented, using this function resulted in a segmentation fault. The removal is consistent with upstream.

Red Hat recommends using the **ssl.SSLContext** class and the **ssl.SSLContext.wrap_socket()** function instead. Most applications can simply use the **ssl.create_default_context()** function, which creates a context with secure default settings. The default context uses the system's default trust store, too.

Symbols from libraries linked as dependencies no longer resolved by ld

Previously, the **ld** linker resolved any symbols present in any linked library, even if some libraries were linked only implicitly as dependencies of other libraries. This allowed developers to use symbols from the implicitly linked libraries in application code and omit explicitly specifying these libraries for linking.

For security reasons, **ld** has been changed to not resolve references to symbols in libraries linked implicitly as dependencies.

As a result, linking with **ld** fails when application code attempts to use symbols from libraries not declared for linking and linked only implicitly as dependencies. To use symbols from libraries linked as dependencies, developers must explicitly link against these libraries as well.

To restore the previous behavior of **ld**, use the **-copy-dt-needed-entries** command-line option. (BZ#[1292230](#))

Windows guest virtual machine support limited

As of Red Hat Enterprise Linux 7, Windows guest virtual machines are supported only under specific subscription programs, such as Advanced Mission Critical (AMC).

libnetlink is deprecated

The **libnetlink** library contained in the **iproute-devel** package has been deprecated. The user should use the **libnl** and **libmnl** libraries instead.

S3 and S4 power management states for KVM have been deprecated

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 Certificate Server plug-in udnPwDDirAuth is discontinued

The **udnPwDDirAuth** authentication plug-in for the Red Hat Certificate Server was removed in Red Hat Enterprise Linux 7.3. Profiles using the plug-in are no longer supported. Certificates created with a profile using the **udnPwDDirAuth** plug-in are still valid if they have been approved.

Red Hat Access plug-in for IdM is discontinued

The Red Hat Access plug-in for Identity Management (IdM) was removed in Red Hat Enterprise Linux 7.3. During the update, the **redhat-access-plugin-ipa** package is automatically uninstalled. Features previously provided by the plug-in, such as Knowledgebase access and support case engagement, are still available through the Red Hat Customer Portal. Red Hat recommends to explore alternatives, such as the **redhat-support-tool** tool.

The Ipsilon identity provider service for federated single sign-on

The **ipsilon** packages were introduced as Technology Preview in Red Hat Enterprise Linux 7.2. Ipsilon links authentication providers and applications or utilities to allow for single sign-on (SSO).

Red Hat does not plan to upgrade Ipsilon from Technology Preview to a fully supported feature. The **ipsilon** packages will be removed from Red Hat Enterprise Linux in a future minor release.

Red Hat has released Red Hat Single Sign-On as a web SSO solution based on the Keycloak community project. Red Hat Single Sign-On provides greater capabilities than Ipsilon and is designated as the standard web SSO solution across the Red Hat product portfolio.

Several rsyslog options deprecated

The **rsyslog** utility version in Red Hat Enterprise Linux 7.4 has deprecated a large number of options. These options no longer have any effect and cause a warning to be displayed.

- The functionality previously provided by the options **-c**, **-u**, **-q**, **-x**, **-A**, **-Q**, **-4**, and **-6** can be achieved using the **rsyslog** configuration.
- There is no replacement for the functionality previously provided by the options **-l** and **-s**

Deprecated symbols from the memkind library

The following symbols from the **memkind** library have been deprecated:

- **memkind_finalize()**
- **memkind_get_num_kind()**
- **memkind_get_kind_by_partition()**
- **memkind_get_kind_by_name()**
- **memkind_partition_mmap()**
- **memkind_get_size()**
- **MEMKIND_ERROR_MEMALIGN**
- **MEMKIND_ERROR_MALLCTL**
- **MEMKIND_ERROR_GETCPU**
- **MEMKIND_ERROR_PMTT**
- **MEMKIND_ERROR_TIEDISTANCE**
- **MEMKIND_ERROR_ALIGNMENT**
- **MEMKIND_ERROR_MALLOCX**
- **MEMKIND_ERROR_REPNAME**
- **MEMKIND_ERROR_PTHREAD**
- **MEMKIND_ERROR_BADPOLICY**
- **MEMKIND_ERROR_REPPOLICY**

Options of Sockets API Extensions for SCTP (RFC 6458) deprecated

The options **SCTP_SNDRCV**, **SCTP_EXTRCV** and **SCTP_DEFAULT_SEND_PARAM** of Sockets API Extensions for the Stream Control Transmission Protocol have been deprecated per the RFC 6458 specification.

New options **SCTP_SNDINFO**, **SCTP_NXTINFO**, **SCTP_NXTINFO** and **SCTP_DEFAULT_SNDINFO** have been implemented as a replacement for the deprecated options.

Managing NetApp ONTAP using SSLv2 and SSLv3 is no longer supported by **libstorageMgmt**

The SSLv2 and SSLv3 connections to the NetApp ONTAP storage array are no longer supported by the **libstorageMgmt** library. Users can contact NetApp support to enable the Transport Layer Security (TLS) protocol.

dconf-dbus-1 has been deprecated and **dconf-editor** is now delivered separately

With this update, the **dconf-dbus-1** API has been removed. However, the **dconf-dbus-1** library has been backported to preserve binary compatibility. Red Hat recommends using the **GDBus** library instead of **dconf-dbus-1**.

The **dconf-error.h** file has been renamed to **dconf-enums.h**. In addition, the **dconf Editor** is now delivered in the separate dconf-editor package.

FreeRADIUS no longer accepts `Auth-Type := System`

The **FreeRADIUS** server no longer accepts the **`Auth-Type := System`** option for the **`rlm_unix`** authentication module. This option has been replaced by the use of the **`unix`** module in the **`authorize`** section of the configuration file.

Deprecated Device Drivers

The following device drivers continue to be supported until the end of life of Red Hat Enterprise Linux 7 but will likely not be supported in future major releases of this product and are not recommended for new deployments.

- 3w-9xxx
- 3w-sas
- aic79xx
- aoe
- arcmsr
- ata drivers:
 - acard-ahci
 - sata_mv
 - sata_nv
 - sata_promise
 - sata_qstor
 - sata_sil
 - sata_sil24
 - sata_sis
 - sata_svw
 - sata_sx4
 - sata_uli
 - sata_via
 - sata_vsc
- bfa
- cxgb3
- cxgb3i

- hptiop
- initio
- isci
- iw_cxgb3
- mptbase
- mptctl
- mptsas
- mptscsih
- mptspi
- mtip32xx
- mvsas
- mvumi
- OSD drivers:
 - osd
 - libosd
- osst
- pata drivers:
 - pata_acpi
 - pata_ali
 - pata_amd
 - pata_arasan_cf
 - pata_artop
 - pata_atiixp
 - pata_atp867x
 - pata_cmd64x
 - pata_cs5536
 - pata_hpt366
 - pata_hpt37x
 - pata_hpt3x2n

- pata_hpt3x3
- pata_it8213
- pata_it821x
- pata_jmicron
- pata_marvell
- pata_netcell
- pata_ninja32
- pata_oldpiix
- pata_pdc2027x
- pata_pdc202xx_old
- pata_piccolo
- pata_rdc
- pata_sch
- pata_serverworks
- pata_sil680
- pata_sis
- pata_via
- pdc_adma
- pm80xx(pm8001)
- pmcraid
- qla3xxx
- stex
- sx8
- ufshcd
- wireless drivers:
 - carl9170
 - iwl4965
 - iwl3945
 - mwl8k

- rt73usb
- rt61pci
- rtl8187
- wil6210

Deprecated Adapters

- The following adapters from the **aacraid** driver have been deprecated:
 - PERC 2/Si (Iguana/PERC2Si), PCI ID 0x1028:0x0001
 - PERC 3/Di (Opal/PERC3Di), PCI ID 0x1028:0x0002
 - PERC 3/Si (SlimFast/PERC3Si), PCI ID 0x1028:0x0003
 - PERC 3/Di (Iguana FlipChip/PERC3DiF), PCI ID 0x1028:0x0004
 - PERC 3/Di (Viper/PERC3DiV), PCI ID 0x1028:0x0002
 - PERC 3/Di (Lexus/PERC3DiL), PCI ID 0x1028:0x0002
 - PERC 3/Di (Jaguar/PERC3DiJ), PCI ID 0x1028:0x000a
 - PERC 3/Di (Dagger/PERC3DiD), PCI ID 0x1028:0x000a
 - PERC 3/Di (Boxster/PERC3DiB), PCI ID 0x1028:0x000a
 - catapult, PCI ID 0x9005:0x0283
 - tomcat, PCI ID 0x9005:0x0284
 - Adaptec 2120S (Crusader), PCI ID 0x9005:0x0285
 - Adaptec 2200S (Vulcan), PCI ID 0x9005:0x0285
 - Adaptec 2200S (Vulcan-2m), PCI ID 0x9005:0x0285
 - Legend S220 (Legend Crusader), PCI ID 0x9005:0x0285
 - Legend S230 (Legend Vulcan), PCI ID 0x9005:0x0285
 - Adaptec 3230S (Harrier), PCI ID 0x9005:0x0285
 - Adaptec 3240S (Tornado), PCI ID 0x9005:0x0285
 - ASR-2020ZCR SCSI PCI-X ZCR (Skyhawk), PCI ID 0x9005:0x0285
 - ASR-2025ZCR SCSI SO-DIMM PCI-X ZCR (Terminator), PCI ID 0x9005:0x0285
 - ASR-2230S + ASR-2230SLP PCI-X (Lancer), PCI ID 0x9005:0x0286
 - ASR-2130S (Lancer), PCI ID 0x9005:0x0286
 - AAR-2820SA (Intruder), PCI ID 0x9005:0x0286

- AAR-2620SA (Intruder), PCI ID 0x9005:0x0286
- AAR-2420SA (Intruder), PCI ID 0x9005:0x0286
- ICP9024RO (Lancer), PCI ID 0x9005:0x0286
- ICP9014RO (Lancer), PCI ID 0x9005:0x0286
- ICP9047MA (Lancer), PCI ID 0x9005:0x0286
- ICP9087MA (Lancer), PCI ID 0x9005:0x0286
- ICP5445AU (Hurricane44), PCI ID 0x9005:0x0286
- ICP9085LI (Marauder-X), PCI ID 0x9005:0x0285
- ICP5085BR (Marauder-E), PCI ID 0x9005:0x0285
- ICP9067MA (Intruder-6), PCI ID 0x9005:0x0286
- Themisto Jupiter Platform, PCI ID 0x9005:0x0287
- Themisto Jupiter Platform, PCI ID 0x9005:0x0200
- Callisto Jupiter Platform, PCI ID 0x9005:0x0286
- ASR-2020SA SATA PCI-X ZCR (Skyhawk), PCI ID 0x9005:0x0285
- ASR-2025SA SATA SO-DIMM PCI-X ZCR (Terminator), PCI ID 0x9005:0x0285
- AAR-2410SA PCI SATA 4ch (Jaguar II), PCI ID 0x9005:0x0285
- CERC SATA RAID 2 PCI SATA 6ch (DellCorsair), PCI ID 0x9005:0x0285
- AAR-2810SA PCI SATA 8ch (Corsair-8), PCI ID 0x9005:0x0285
- AAR-21610SA PCI SATA 16ch (Corsair-16), PCI ID 0x9005:0x0285
- ESD SO-DIMM PCI-X SATA ZCR (Prowler), PCI ID 0x9005:0x0285
- AAR-2610SA PCI SATA 6ch, PCI ID 0x9005:0x0285
- ASR-2240S (SabreExpress), PCI ID 0x9005:0x0285
- ASR-4005, PCI ID 0x9005:0x0285
- IBM 8i (AvonPark), PCI ID 0x9005:0x0285
- IBM 8i (AvonPark Lite), PCI ID 0x9005:0x0285
- IBM 8k/8k-l8 (Aurora), PCI ID 0x9005:0x0286
- IBM 8k/8k-l4 (Aurora Lite), PCI ID 0x9005:0x0286
- ASR-4000 (BlackBird), PCI ID 0x9005:0x0285
- ASR-4800SAS (Marauder-X), PCI ID 0x9005:0x0285

- ASR-4805SAS (Marauder-E), PCI ID 0x9005:0x0285
- ASR-3800 (Hurricane44), PCI ID 0x9005:0x0286
- Perc 320/DC, PCI ID 0x9005:0x0285
- Adaptec 5400S (Mustang), PCI ID 0x1011:0x0046
- Adaptec 5400S (Mustang), PCI ID 0x1011:0x0046
- Dell PERC2/QC, PCI ID 0x1011:0x0046
- HP NetRAID-4M, PCI ID 0x1011:0x0046
- Dell Catchall, PCI ID 0x9005:0x0285
- Legend Catchall, PCI ID 0x9005:0x0285
- Adaptec Catch All, PCI ID 0x9005:0x0285
- Adaptec Rocket Catch All, PCI ID 0x9005:0x0286
- Adaptec NEMER/ARK Catch All, PCI ID 0x9005:0x0288
- The following adapters from the **mpt2sas** driver have been deprecated:
 - SAS2004, PCI ID 0x1000:0x0070
 - SAS2008, PCI ID 0x1000:0x0072
 - SAS2108_1, PCI ID 0x1000:0x0074
 - SAS2108_2, PCI ID 0x1000:0x0076
 - SAS2108_3, PCI ID 0x1000:0x0077
 - SAS2116_1, PCI ID 0x1000:0x0064
 - SAS2116_2, PCI ID 0x1000:0x0065
 - SSS6200, PCI ID 0x1000:0x007E
- The following adapters from the **megaraid_sas** driver have been deprecated:
 - Dell PERC5, PCI ID 0x1028:0x15
 - SAS1078R, PCI ID 0x1000:0x60
 - SAS1078DE, PCI ID 0x1000:0x7C
 - SAS1064R, PCI ID 0x1000:0x411
 - VERDE_ZCR, PCI ID 0x1000:0x413
 - SAS1078GEN2, PCI ID 0x1000:0x78
 - SAS0079GEN2, PCI ID 0x1000:0x79

- SAS0073SKINNY, PCI ID 0x1000:0x73
- SAS0071SKINNY, PCI ID 0x1000:0x71
- The following adapters from the **qla2xxx** driver have been deprecated:
 - ISP24xx, PCI ID 0x1077:0x2422
 - ISP24xx, PCI ID 0x1077:0x2432
 - ISP2422, PCI ID 0x1077:0x5422
 - QLE220, PCI ID 0x1077:0x5432
 - QLE81xx, PCI ID 0x1077:0x8001
 - QLE10000, PCI ID 0x1077:0xF000
 - QLE84xx, PCI ID 0x1077:0x8044
 - QLE8000, PCI ID 0x1077:0x8432
 - QLE82xx, PCI ID 0x1077:0x8021
- The following adapters from the **qla4xxx** driver have been deprecated:
 - QLOGIC_ISP8022, PCI ID 0x1077:0x8022
 - QLOGIC_ISP8324, PCI ID 0x1077:0x8032
 - QLOGIC_ISP8042, PCI ID 0x1077:0x8042
- The following adapters from the **be2iscsi** driver have been deprecated:
 - BladeEngine 2 (BE2) Devices
 - BladeEngine2 10Gb iSCSI Initiator (generic), PCI ID 0x19a2:0x212
 - OneConnect OCe10101, OCm10101, OCe10102, OCm10102 BE2 adapter family, PCI ID 0x19a2:0x702
 - OCe10100 BE2 adapter family, PCI ID 0x19a2:0x703
 - BladeEngine 3 (BE3) Devices
 - OneConnect TOMCAT iSCSI, PCI ID 0x19a2:0x0712
 - BladeEngine3 iSCSI, PCI ID 0x19a2:0x0222
- The following Ethernet adapters controlled by the **be2net** driver have been deprecated:
 - BladeEngine 2 (BE2) Devices
 - OneConnect TIGERSHARK NIC, PCI ID 0x19a2:0700
 - BladeEngine2 Network Adapter, PCI ID 0x19a2:0211
 - BladeEngine 3 (BE3) Devices

- OneConnect TOMCAT NIC, PCI ID 0x19a2:0x0710
- BladeEngine3 Network Adapter, PCI ID 0x19a2:0221
- The following adapters from the **lpfc** driver have been deprecated:
 - BladeEngine 2 (BE2) Devices
 - OneConnect TIGERSHARK FCoE, PCI ID 0x19a2:0x0704
 - BladeEngine 3 (BE3) Devices
 - OneConnect TOMCAT FCoE, PCI ID 0x19a2:0x0714
 - Fibre Channel (FC) Devices
 - FIREFLY, PCI ID 0x10df:0x1ae5
 - PROTEUS_VF, PCI ID 0x10df:0xe100
 - BALIUS, PCI ID 0x10df:0xe131
 - PROTEUS_PF, PCI ID 0x10df:0xe180
 - RFLY, PCI ID 0x10df:0xf095
 - PFLY, PCI ID 0x10df:0xf098
 - LP101, PCI ID 0x10df:0xf0a1
 - TFLY, PCI ID 0x10df:0xf0a5
 - BSMB, PCI ID 0x10df:0xf0d1
 - BMID, PCI ID 0x10df:0xf0d5
 - ZSMB, PCI ID 0x10df:0xf0e1
 - ZMID, PCI ID 0x10df:0xf0e5
 - NEPTUNE, PCI ID 0x10df:0xf0f5
 - NEPTUNE_SCSP, PCI ID 0x10df:0xf0f6
 - NEPTUNE_DCSP, PCI ID 0x10df:0xf0f7
 - FALCON, PCI ID 0x10df:0xf180
 - SUPERFLY, PCI ID 0x10df:0xf700
 - DRAGONFLY, PCI ID 0x10df:0xf800
 - CENTAUR, PCI ID 0x10df:0xf900
 - PEGASUS, PCI ID 0x10df:0xf980
 - THOR, PCI ID 0x10df:0xfa00

- VIPER, PCI ID 0x10df:0xfb00
- LP10000S, PCI ID 0x10df:0xfc00
- LP11000S, PCI ID 0x10df:0xfc10
- LPE11000S, PCI ID 0x10df:0xfc20
- PROTEUS_S, PCI ID 0x10df:0xfc50
- HELIOS, PCI ID 0x10df:0xfd00
- HELIOS_SCSP, PCI ID 0x10df:0xfd11
- HELIOS_DCSP, PCI ID 0x10df:0xfd12
- ZEPHYR, PCI ID 0x10df:0xfe00
- HORNET, PCI ID 0x10df:0xfe05
- ZEPHYR_SCSP, PCI ID 0x10df:0xfe11
- ZEPHYR_DCSP, PCI ID 0x10df:0xfe12
- Lancer FCoE CNA Devices
 - OCe15104-FM, PCI ID 0x10df:0xe260
 - OCe15102-FM, PCI ID 0x10df:0xe260
 - OCm15108-F-P, PCI ID 0x10df:0xe260

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

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

The **libcxgb3** library and the **cxgb3** firmware package have been deprecated

The **libcxgb3** library provided by the **libibverbs** package and the **cxgb3** firmware package have been deprecated. They continue to be supported in Red Hat Enterprise Linux 7 but will likely not be supported in the next major releases of this product. This change corresponds with the deprecation of the **cxgb3**, **cxgb3i**, and **iw_cxgb3** drivers listed above.

SFN4XXX adapters have been deprecated

Starting with Red Hat Enterprise Linux 7.4, SFN4XXX Solarflare network adapters have been deprecated. Previously, Solarflare had a single driver **sfc** for all adapters. Recently, support of SFN4XXX was split from **sfc** and moved into a new SFN4XXX-only driver, called **sfc-falcon**. Both drivers continue to be supported at this time, but **sfc-falcon** and SFN4XXX support is scheduled for removal in a future major release.

Software-initiated-only FCoE storage technologies have been deprecated

The software-initiated-only type of the Fibre Channel over Ethernet (FCoE) storage technology has been deprecated due to limited customer adoption. The software-initiated-only storage technology will remain supported for the life of Red Hat Enterprise Linux 7. The deprecation notice indicates the intention to remove software-initiated-based FCoE support in a future major release of Red Hat Enterprise Linux.

It is important to note that the hardware support and the associated user-space tools (such as drivers, **libfc**, or **libfcoe**) are unaffected by this deprecation notice.

Containers using the `libvirt-lxc` tooling have been deprecated

The following `libvirt-lxc` packages are deprecated since Red Hat Enterprise Linux 7.1:

- `libvirt-daemon-driver-lxc`
- `libvirt-daemon-lxc`
- `libvirt-login-shell`

Future development on the Linux containers framework is now based on the **docker** command-line interface. `libvirt-lxc` tooling may be removed in a future release of Red Hat Enterprise Linux (including Red Hat Enterprise Linux 7) and should not be relied upon for developing custom container management applications.

For more information, see the [Red Hat KnowledgeBase article](#).

The Perl and shell scripts for Directory Server have been deprecated

The Perl and shell scripts, which are provided by the `389-ds-base` package, have been deprecated. The scripts will be replaced by new utilities in the next major release of Red Hat Enterprise Linux.

The [Shell Scripts](#) and [Perl Scripts](#) sections in the *Red Hat Directory Server Command, Configuration, and File Reference* have been updated. The descriptions of affected scripts contain now a note that they are deprecated.

`libguestfs` can no longer inspect ISO installer files

The `libguestfs` library does no longer support inspecting ISO installer files, for example using the `guestfish` or `virt-inspector` utilities. Use the `osinfo-detect` command for inspecting ISO files instead. This command can be obtained from the `libosinfo` package.

Creating internal snapshots of virtual machines has been deprecated

Due to their lack of optimization and stability, internal virtual machine snapshots are now deprecated. In their stead, external snapshots are recommended for use. For more information, including instructions for creating external snapshots, see the [Virtualization Deployment and Administration Guide](#).

IVSHMEM has been deprecated

The inter-VM shared memory device (IVSHMEM) feature has been deprecated. Therefore, in a future major release of RHEL, if a virtual machine (VM) is configured to share memory between multiple virtual machines in the form of a PCI device that exposes memory to guests, the VM will fail to boot.

The `gnome-shell-browser-plugin` subpackage has been deprecated

Since the Firefox Extended Support Release (ESR 60), Firefox no longer supports the Netscape Plugin Application Programming Interface (NPAPI) that was used by the `gnome-shell-browser-plugin` subpackage. The subpackage, which provided the functionality to install GNOME Shell Extensions, has thus been deprecated. The installation of GNOME Shell Extensions is now handled directly in the `gnome-software` package.

The VDO read cache has been deprecated

The read cache functionality in Virtual Data Optimizer (VDO) has been deprecated. The read cache is disabled by default on new VDO volumes.

In the next major Red Hat Enterprise Linux release, the read cache functionality will be removed, and you will no longer be able to enable it using the `--readCache` option of the `vdo` utility.

`cpuid` has been deprecated

The **cpuid** command has been deprecated. A future major release of Red Hat Enterprise Linux will no longer support using **cpuid** to dump the information about CPUID instruction for each CPU. To obtain similar information, use the **lscpu** command instead.

KDE has been deprecated

KDE Plasma Workspaces (KDE), which has been provided as an alternative to the default GNOME desktop environment has been deprecated. A future major release of Red Hat Enterprise Linux will no longer support using KDE instead of the default GNOME desktop environment.

The `lwresd` daemon has been deprecated

The **lwresd** daemon, which is a part of the `bind` package, has been deprecated. A future major release of Red Hat Enterprise Linux will no longer support providing name lookup services to clients that use the BIND 9 lightweight resolver library with **lwresd**.

The recommended replacements are:

- The **systemd-resolved** daemon and **nss-resolve** API, provided by the `systemd` package
- The **unbound** library API and daemon, provided by the `unbound` and `unbound-libs` packages
- The **getaddrinfo** glibc library call and related

PART V. KNOWN ISSUES

This part documents known problems in Red Hat Enterprise Linux 7.6.

CHAPTER 52. AUTHENTICATION AND INTEROPERABILITY

RADIUS proxy functionality is now also available in IdM running in FIPS mode

In FIPS mode, OpenSSL disables the use of the MD5 digest algorithm by default. Consequently, because the RADIUS protocol requires MD5 to encrypt a secret between the RADIUS client and the RADIUS server, the unavailability of MD5 in FIPS mode causes the RHEL Identity Management (IdM) RADIUS proxy server to fail.

If the RADIUS server is running on the same host as the IdM master, you can work around the problem and enable MD5 within the secure perimeter.

To do that, create a file `/etc/systemd/system/radiusd.service.d/ipa-otp.conf` with the following content:

```
# /etc/systemd/system/radiusd.service.d/ipa-otp.conf
[Service]
Environment=OPENSSL_FIPS_NON_APPROVED_MD5_ALLOW=1
```

To apply the change, reload the **systemd** configuration:

```
# systemctl daemon-reload
```

and start the **radiusd** service:

```
# systemctl start radiusd
```

The configuration of the RADIUS proxy requires the use of a common secret between the client and the server to wrap credentials. Specify this secret in the configuration of the RADIUS proxy in RHEL IdM using the command line interface (CLI) or web UI. To do it in the CLI:

```
# ipa radiusproxy-add name_of_your_proxy_server --secret your_secret
```

(BZ#[1571754](#))

CHAPTER 53. COMPILER AND TOOLS

GCC thread sanitizer included in RHEL no longer works

Due to incompatible changes in kernel memory mapping, the thread sanitizer included with the GNU C Compiler (GCC) compiler version in RHEL no longer works. Additionally, the thread sanitizer cannot be adapted to the incompatible memory layout. As a result, it is no longer possible to use the GCC thread sanitizer included with RHEL.

As a workaround, use the version of GCC included in Red Hat Developer Toolset to build code which uses the thread sanitizer. (BZ#1569484)

CHAPTER 54. DESKTOP

Firefox 60.1 ESR fails to start on IBM Z and POWER

JavaScript engine in the **Firefox** 60.1 Extended Support Release (ESR) browser was changed. As a consequence, **Firefox** 60.1 ESR on IBM Z and POWER architectures fails to start with a segmentation fault error message. (BZ#[1576289](#), BZ#1579705)

GV100GL graphics cannot use correctly more than one monitor

Due to missing signed firmware for the GV100GL graphics, GV100GL cannot have more than one monitor connected. When a second monitor is connected, it is recognized, and graphics set the correct resolution, but the monitor stays in power-saving mode. To work around this problem, install the **NVIDIA** binary driver. As a result, the second monitor output works as expected under the described circumstances. (BZ#[1624337](#))

The Files application can not burn disks in default installation

The default installation of the **Files** application does not include the **brasero-nautilus** package necessary for burning CDs or DVDs. As a consequence, the **Files** application allows files to be dragged and dropped into CD or DVD devices but no content is burned to the CD or DVD. As a workaround, install **brasero-nautilus** package by:

```
# yum install brasero-nautilus
```

(BZ#[1600163](#))

The on screen keyboard feature not visible in GTK applications

After enabling the **on screen keyboard** feature by using the **Settings - Universal Access - Typing - Screen keyboard** menu, **on screen keyboard** is not visible to access with GIMP Toolkit (GTK) applications, such as **gedit**.

To work around this problem, add the below line into the **/etc/environment** configuration file, and restart GNOME:

```
GTK_IM_MODULE=ibus
```

(BZ#[1625700](#))

32- and 64-bit fwupd packages cannot be used together when installing or upgrading the system

The **/usr/lib/systemd/system/fwupd.service** file in the **fwupd** packages is different for 32- and 64-bit architectures. Consequently, it is impossible to install both 32- and 64-bit **fwupd** packages or to upgrade a Red Hat Enterprise Linux 7.5 system with both 32- and 64-bit **fwupd** packages to Red Hat Enterprise Linux 7.6.

To work around this problem:

- Either do not install multilibary **fwupd** packages.
- Or remove the 32-bit or the 64-bit **fwupd** package before upgrading from Red Hat Enterprise Linux 7.5 to Red Hat Enterprise Linux 7.6. (BZ#[1623466](#))

Installation in and booting into graphical mode are not possible on Huawei servers

When installing RHEL 7.6 in graphical mode on Huawei servers with AMD64 and Intel 64 processors, the screen becomes blurred and the install interface is no longer visible. After finishing the installation in console mode, the operating system cannot be booted into graphical mode.

To work around this problem:

1. Add kernel command line parameter **inst.xdriver=fbdev** when installing the system, and install the system as **server with GUI**. 2. After the installation completes, reboot and add kernel command line **single** to make the system boot into maintenance mode. 3. Run the following commands:

```
rpm -e xorg-x11-drivers
rpm -e xorg-x11-drv-vesa
init 5
```

(BZ#[1624847](#))

X.org server crashes during fast user switching

The X.Org X11 **qx1** video driver does not emulate the leaving virtual terminal event on shutdown. Consequently, the X.Org display server terminates unexpectedly during fast user switching, and the current user session is terminated when switching a user. (BZ#[1640918](#))

X.org X11 crashes on Lenovo T580

Due to a bug in the **libpciaccess** library, the X.org X11 server terminates unexpectedly on Lenovo T580 laptops. (BZ#[1641044](#))

Soft lock-ups might occur during boot in the kernel with **i915**

On a rare occasion when a GM45 system has an improper firmware configuration, an incorrect **DisplayPort** hot-plug signal can cause the **i915** driver to be overloaded on boot. Consequently, certain GM45 systems might experience very slow boot times while the video driver attempts to work around the problem. In some cases, the kernel might report soft lock-ups occurring. Customers are advised to contact their hardware vendors and request a firmware update to address this problem. (BZ#1608704)

System boots to a blank screen when Xinerama is enabled

When the Xinerama extension is enabled in **/etc/X11/xorg.conf** on a system using the **nvidia/nouveau** driver, the RANDR X extension gets disabled. Consequently, login screen fails to start upon boot due to the RANDR X extension being disabled. To work around this problem, do not enable Xinerama in **/etc/X11/xorg.conf**. (BZ#[1579257](#))

CHAPTER 55. FILE SYSTEMS

Mounting a non-existent NFS export outputs a different error than in RHEL 6

The **mount** utility prints the **operation not permitted** error message when an NFS client is trying to mount a server export that does not exist. In Red Hat Enterprise Linux 6, the **access denied** message was printed in the same situation. (BZ#1428549)

XFS disables per-inode DAX functionality

Per-inode direct access (DAX) options are now disabled in the XFS file system due to unresolved issues with this feature. XFS now ignores existing per-inode DAX flags on the disk.

You can still set file system DAX behavior using the **dax** mount option:

```
# mount -o dax device mount-point
```

(BZ#1623150)

CHAPTER 56. INSTALLATION AND BOOTING

Certain RPM packages are not available on binary DVDs

The virt-p2v RPM, syslinux-tftpbboot, LibreOffice, and KDE language packages are not available on Red Hat Enterprise Linux binary DVDs due to the limited size of a single-layered DVD. The packages are still consumable by Red Hat Subscription Management and Red Hat Network by enabling the relevant updates after Anaconda installation. The packages can also be downloaded from <https://access.redhat.com/downloads>.

In addition, the syslinux-tftpbboot package has been moved from the Optional channel to the Base channel (Server variant) and it is now also available for the IBM POWER, little endian architecture. (BZ#[1611665](#), BZ#1592748, BZ#1616396)

The content location detection code is not working on Red Hat Virtualization Hosts

Red Hat Virtualization Hosts cannot select the hardening profile from locally-installed content. To work around this problem, use the oscap-anaconda-addon package to fetch the Red Hat Enterprise Linux datastream file from a URL.

1. Upload the **ssg-rhel7-ds.xml** datastream file from the Red Hat Enterprise Linux 7 scap-security-guide package to your network so it can be discovered by Anaconda.

To do so:

a) Use Python to set up a web server in a directory that contains the **ssg-rhel7-ds.xml** datastream file and listens on port 8000. Example: `python2 -m SimpleHTTPServer`, or `python3 -m http.server`.

or,

b) Upload the **ssg-rhel7-ds.xml** datastream file to a HTTPS or FTP Server.

2. In the **Security Policy** window of Anaconda's Graphical User Interface, click **Change Content** and enter the URL that points to the **ssg-rhel7-ds.xml** datastream file, for example: `http://gateway:8000/ssg-rhel7-ds.xml` or `ftp://my-ftp-server/ssg-rhel7-ds.xml`.

The **ssg-rhel7-ds.xml** datastream file is now available and Red Hat Virtualization Hosts can select the hardening profile. (BZ#[1636847](#))

Composer can not create live ISO system images

Because of a dependency resolution issue, tools for building live ISO images are not recognized by Composer at run time. As a consequence, Composer fails to build live ISO images and users can not create this type of system image.

Composer is available as a Technology Preview. (BZ#1642156)

CHAPTER 57. KERNEL

Cache information is missing in `sysfs` if firmware does not support ACPI PPTT

The kernel-alt package has been updated to use the Advanced Configuration and Power Interface Processor Properties Topology Table (ACPI PPTT) to populate CPU topology including the CPU's cache information. Consequently, on systems whose firmware does not support ACPI PPTT, the `/sys/devices/system/cpu/cpu0/cache` file does not contain the cache information. To work around this problem, check for updated firmware that includes ACPI PPTT support with your hardware vendor. (BZ#1615370)

PCI-passthrough of devices connected to PCIe slots is not possible with default settings of HPE ProLiant Gen8 and Gen9

Default settings of HPE ProLiant Gen8 and Gen9 systems disallow use of PCI-passthrough for devices connected to PCIe slots. Consequently, any attempt to pass through such devices fails with the following message in the kernel log:

```
Device is ineligible for IOMMU domain attach due to platform RMRR requirement. Contact your platform vendor.
```

To work around this problem:

- In case of HPE ProLiant Gen8, reconfigure mentioned system settings with the **conrep** tool provided by HPE.
- In case of HPE ProLiant Gen9, update system firmware or NICs firmware depending on type of used NICs.

For more details about the workaround, see https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-c04781229. (BZ#1615210)

Attaching a non-RoCE device to RXE driver no longer causes a kernel to panic

When a user created a Soft RDMA Over Converged Ethernet (Soft RoCE) interface and attached a non-RoCE device, certain issues were observed in the RXE driver. As a consequence, a kernel panicked when rebooting or shutting down a host. With this update, disabling the Soft RoCE interface before rebooting or shutting down a host fixes the issue. As a result, the host no longer panics in the described scenario. (BZ#1520302)

Enabling the `bcc` packages for the 64-bit AMD and Intel architectures only

The BPF Compiler Collection (BCC) library and the **pcp-pmda-bcc** plugins use the **bpf()** system call, which is enabled only on the 64-bit AMD and Intel CPU architectures. As a result, Red Hat Enterprise Linux 7 only supports **BCC** and **pcp-pmda-bcc** for the 64-bit AMD and Intel CPU architectures. (BZ#1633185)

Kernel panics can occur on virtual machines that use SEV

Currently, the Secure Encrypted Virtualization (SEV) feature does not work correctly under high virtio traffic. As a consequence, a kernel panic may in some cases occur on a virtual machine that uses SEV, which generates the **DMA: Out of SW-IOMMU space** error message.

To work around this problem, reserve 512 MiB for the Software Input/Output Translation Lookaside Buffer (SWIOTLB) by appending the **swiotlb=262144** parameter to the guest kernel command line. (BZ#1637992)

Branch prediction of ternary operators no longer causes a system panic

Previously, the branch prediction of ternary operators caused that the compiler incorrectly called the

blk_queue_nonrot() function before checking the **mddev->queue** structure. As a consequence, the system panicked. With this update, checking **mddev->queue** and then calling **blk_queue_nonrot()** prevents the bug from appearing. As a result, the system no longer panics in the described scenario. (BZ#1627563)

RAID1 write-behind causes a kernel panic

Write-behind mode in the Redundant Array of Independent Disks Mode 1 (RAID1) virtualization technology uses the upper layer bio structures, which are freed immediately after the bio structures written to bottom layer disks come back. As a consequence, a kernel panic is triggered and the **write-behind** function cannot be used. (BZ#1632575)

CHAPTER 58. NETWORKING

Verification of signatures using the MD5 hash algorithm is disabled in Red Hat Enterprise Linux 7

It is impossible to connect to any Wi-Fi Protected Access (WPA) Enterprise Access Point (AP) that requires MD5 signed certificates. To work around this problem, copy the `wpa_supplicant.service` file from the `/usr/lib/systemd/system/` directory to the `/etc/systemd/system/` directory and add the following line to the Service section of the file:

```
Environment=OPENSSL_ENABLE_MD5_VERIFY=1
```

Then run the **`systemctl daemon-reload`** command as root to reload the service file.

Important: Note that MD5 certificates are highly insecure and Red Hat does not recommend using them. (BZ#1062656)

CHAPTER 59. SECURITY

OpenSCAP rpmverifypackage does not work correctly

The **chdir** and **chroot** system calls are called twice by the **rpmverifypackage** probe. Consequently, an error occurs when the probe is utilized during an **OpenSCAP** scan with custom Open Vulnerability and Assessment Language (OVAL) content.

To work around this problem, do not use the **rpmverifypackage_test** OVAL test in your content or use only the content from the **scap-security-guide** package where **rpmverifypackage_test** is not used. (BZ#[1603347](#))

dconf databases are not checked by OVAL

OVAL (Open Vulnerability and Assessment Language) checks used in the **SCAP Security Guide** project are not able to read a **dconf** binary database, only files used to generate the database. The database is not regenerated automatically, the administrator needs to enter the **dconf update** command. As a consequence, changes to the database that are not reflected in files in the **/etc/dconf/db/** directory cannot be a part of scanning. This might result in false negatives.

To work around this problem, run **dconf update** periodically, for example, using the **/etc/crontab** configuration file. (BZ#[1631378](#))

SCAP Workbench fails to generate results-based remediations from tailored profiles

The following error occurs when trying to generate results-based remediation roles from a customized profile using the **SCAP Workbench** tool:

```
Error generating remediation role '.../remediation.sh': Exit code of
'oscap' was 1: [output truncated]
```

To work around this problem, use the **oscap** command with the **--tailoring-file** option. (BZ#[1533108](#))

OpenSCAP scanner results contain a lot of SELinux context error messages

The **OpenSCAP** scanner logs inability to get SELinux context on the **ERROR** level even in situations where it is not a true error. As a result, **OpenSCAP** scanner results contain a lot of SELinux context error messages. Both the **oscap** command-line utility and the **SCAP Workbench** graphical utility outputs can be hard to read for that reason. (BZ#[1640522](#))

CHAPTER 60. SERVERS AND SERVICES

Rsyslog cannot proceed if the default maximum of open files is exceeded

Rsyslog sometimes runs over the default limits for maximum number of open files. Consequently, **rsyslog** cannot open new files.

To work around this problem, modify the rsyslog configuration by increasing this limit to align with systemd-journald. To do so, create a drop-in file named `/etc/systemd/system/rsyslog.service.d/increase_nofile_limit.conf` with the following content:

```
[Service]
LimitNOFILE=16384
```

(BZ#[1553700](#))

Tuned does not set kernel boot command line parameters

The **Tuned** tool does not support Boot Loader Specification (BLS), which is enabled by default. Consequently, **Tuned** does not set certain kernel boot command line parameters, which causes some issues, such as performance decrease or CPU cores not being isolated. To work around this problem, disable BLS and restart **Tuned**.

1. Install the grubby package.
2. Remove the following line from the `/etc/default/grub` file:

```
GRUB_ENABLE_BLSCFG=true
```

3. Re-generate the **grub2.cfg** file by running for non-EFI systems:

```
grub2-mkconfig -o /etc/grub2.cfg
```

or for EFI systems:

```
grub2-mkconfig -o /etc/grub2-efi.cfg
```

4. Restart **Tuned** by running:

```
systemctl restart tuned
```

As a result, **Tuned** sets the kernel boot parameters as expected. (BZ#[1576435](#))

CHAPTER 61. STORAGE

LVM does not support event-based autoactivation of incomplete volume groups

If a volume group is not complete and physical volumes are missing, LVM does not support automatic LVM event-based activation of that volume group. This implies a setting of **--activationmode complete** whenever autoactivation takes place. For information on the **--activationmode complete** option and automatic activation, see the **vgchange(8)** and **pvscan(8)** man pages.

Note that the event-driven autoactivation hooks are enabled when **lvmetad** is enabled with the **global/use_lvmetad=1** setting in the **/etc/lvm/lvm.conf** configuration file. Also note that without autoactivation, there is a direct activation hook at the exact time during boot at which the volume groups are activated with only the physical volumes that are available at that time. Any physical volumes that appear later are not taken into account.

This issue does not affect early boot in **initramfs (dracut)** nor does this affect direct activation from the command line using **vgchange** and **lvchange** calls, which default to **degraded** activation mode. (BZ#1337220)

The vdo service is disabled after upgrading to Red Hat Enterprise Linux 7.6

Upgrading from Red Hat Enterprise Linux 7.5 to 7.6 disables the **vdo** service if it was previously enabled. This is because of missing **systemd** macros in the **vdo** RPM package.

The problem has been fixed in the 7.6 release, and upgrading from Red Hat Enterprise Linux 7.6 to a later release will no longer disable **vdo**. (BZ#1617896)

Data corruption occurs on RAID 10 reshape on top of VDO.

RAID 10 reshape (with both LVM and **mdadm**) on top of VDO corrupts data. Stacking RAID 10 (or other RAID types) on top of VDO does not take advantage of the deduplication and compression capabilities of VDO and is not recommended. (BZ#1528466, BZ#1530776)

System boot is sometimes delayed by ndctl

A **udev** rule installed by the **ndctl** package sometimes delays the system boot process for several minutes on systems with Non-Volatile Dual In-line Memory Module (NVDIMM) devices. In such cases, **systemd** displays a message similar to the following:

```
INFO: task systemd-udevd:1554 blocked for more than 120 seconds.
...
nvdimmm_bus_check_dimm_count+0x31/0xa0 [libnvdimmm]
...
```

To work around the issue, disable the **udev** rule using the following command:

```
# rm /usr/lib/udev/rules.d/80-ndctl.rules
```

After disabling the **udev** rule, the described problem no longer occurs. (BZ#1635441)

LVM might cause data corruption in the first 128kB of allocatable space of a physical volume

A bug in the I/O layer of LVM causes LVM to read and write back the first 128kB of data that immediately follows the LVM metadata on the disk. If another program or the file system is modifying these blocks when you use an LVM command, changes might be lost. As a consequence, this might lead to data corruption in rare cases.

To work around this problem, avoid using LVM commands that change volume group (VG) metadata, such as **lvcreate** or **lvextend**, while logical volumes (LVs) in the VG are in use. (BZ# [1643651](#))

CHAPTER 62. SYSTEM AND SUBSCRIPTION MANAGEMENT

Red Hat Satellite 5.8 availability of RHEL 7.6 EUS, AUS, TUS, and E4S streams delayed

Red Hat Satellite 5 content ISOs are made available on a monthly cadence. Based on this cadence, content ISOs are not available through Red Hat Satellite 5.8 for the following RHEL 7.6 streams at the time of the RHEL 7.6 general availability:

- Extended Update Support (EUS)
- Advanced Update Support (AUS)
- Telco Extended Update Support (TUS)
- Update Services for SAP Solutions (E4S)

The expected delay is two to four weeks. Note that Red Hat Satellite 6 instances are unaffected.

See <https://access.redhat.com/solutions/3621151> for more details. (BZ#1635135)

APPENDIX A. COMPONENT VERSIONS

This appendix provides a list of key components and their versions in the Red Hat Enterprise Linux 7.6 release.

Table A.1. Component Versions

Component	Version
kernel	3.10.0-957
kernel-alt	4.14.0-115
QLogic qla2xxx driver	10.00.00.06.07.6-k
QLogic qla4xxx driver	5.04.00.00.07.02-k0
Emulex lpfc driver	0:12.0.0.5
iSCSI initiator utils (iscsi-initiator-utils)	6.2.0.874-10
DM-Multipath (device-mapper-multipath)	0.4.9-123
LVM (lvm2)	2.02.180-8
qemu-kvm ^[a]	1.5.3-160
qemu-kvm-ma ^[b]	2.12.0-18
<p>[a] The qemu-kvm packages provide KVM virtualization on AMD64 and Intel 64 systems.</p> <p>[b] The qemu-kvm-ma packages provide KVM virtualization on IBM POWER8, IBM POWER9, and IBM Z. Note that KVM virtualization on IBM POWER9 and IBM Z also requires using the kernel-alt packages.</p>	

APPENDIX B. LIST OF BUGZILLAS BY COMPONENT

This appendix provides a list of all components and their related Bugzillas that are included in this book.

Table B.1. List of Bugzillas by Component

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
389-ds-base	BZ# 1560653	BZ# 1515190 , BZ# 1525256 , BZ# 1551071 , BZ# 1552698 , BZ# 1559945 , BZ# 1566444 , BZ# 1568462 , BZ# 1570033 , BZ# 1570649 , BZ# 1576485 , BZ# 1581737 , BZ# 1582092 , BZ# 1582747 , BZ# 1593807 , BZ# 1598478 , BZ# 1598718 , BZ# 1614501		
NetworkManager	BZ# 1414093 , BZ# 1487477	BZ# 1507864		
OVMF			BZ# 653382	
anaconda	BZ# 1562301	BZ# 1360223 , BZ# 1436304 , BZ# 1535781 , BZ# 1554271 , BZ# 1557485 , BZ# 1561662 , BZ# 1561930		
audit	BZ# 1559032			
augeas		BZ# 1544520		
bind	BZ# 1452091 , BZ# 1510008			
binutils		BZ# 1553842 , BZ# 1557346		
clevis	BZ# 1472435			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
cockpit	BZ#1568728			
corosync			BZ#1413573	
criu			BZ#1400230	
custodia			BZ#1403214	
device-mapper-multipath	BZ#1541116, BZ#1554516, BZ#1593459	BZ#1498724, BZ#1526876, BZ#1544958, BZ#1584228, BZ#1610263		
distribution				BZ#1062656
dnf			BZ#1461652	
fence-agents			BZ#1476401	
firefox				BZ#1576289
firewalld	BZ#1477771, BZ#1554993	BZ#1498923		
freetype	BZ#1576504			
fwupd				BZ#1623466
gcc		BZ#1552021		
gcc-libraries	BZ#1600265			
gdb	BZ#1553104	BZ#1347993		
genwqe-tools	BZ#1521050			
ghostscript		BZ#1551782		
git		BZ#1213059, BZ#1284081		
glibc	BZ#1448107, BZ#1461231			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
gnome-shell			BZ#1481395	BZ#1625700
gnutls	BZ#1561481			
ima-evm-utils	BZ#1627278		BZ#1384450	
initscripts	BZ#1493069, BZ#1542514, BZ#1583677	BZ#1554364, BZ#1554690, BZ#1559384, BZ#1572659		
ipa			BZ#1115294, BZ#1298286	
ipa-server-container			BZ#1405325	
ipset	BZ#1557600			
java-11-openjdk	BZ#1570856			
jss	BZ#1557575, BZ#1560682			
kernel	BZ#1205497, BZ#1305092, BZ#1322930, BZ#1344565, BZ#1350553, BZ#1361286, BZ#1451438, BZ#1457161, BZ#1471950, BZ#1496859, BZ#1507027, BZ#1511351, BZ#1515584, BZ#1520356, BZ#1557599, BZ#1570090, BZ#1584753, BZ#1620372	BZ#1527799, BZ#1541250, BZ#1544920, BZ#1554907, BZ#1636930	BZ#916382, BZ#1109348, BZ#1111712, BZ#1206277, BZ#1230959, BZ#1274459, BZ#1299662, BZ#1348508, BZ#1387768, BZ#1393375, BZ#1414957, BZ#1457533, BZ#1460849, BZ#1503123, BZ#1519746, BZ#1589397	BZ#1428549, BZ#1520302, BZ#1528466, BZ#1608704, BZ#1615210, BZ#1623150, BZ#1627563, BZ#1632575, BZ#1637992
kernel-alt				BZ#1615370
kernel-rt	BZ#1297061, BZ#1553351	BZ#1608672		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
kexec-tools	BZ# 1352763			
libcgroup		BZ# 1549175		
libguestfs	BZ# 1541908 , BZ# 1557273		BZ# 1387213 , BZ# 1441197 , BZ# 1477912	
libnftnl	BZ# 1332585			
libpciaccess				BZ# 1641044
libreswan	BZ# 1536404 , BZ# 1591817		BZ# 1375750	
libsepol	BZ# 1564775			
libstoragegmt			BZ# 1119909	
libusnic_verbs			BZ# 916384	
libvirt	BZ# 1447169		BZ# 1283251 , BZ# 1475770	
linuxptp	BZ# 1549015			
llvm-private		BZ# 1417663		
lorax-composer				BZ# 1642156
lvm2				BZ# 1337220 , BZ# 1643651
man-db		BZ# 1515352		
mutter				BZ# 1579257
nautilus				BZ# 1600163
ndctl				BZ# 1635441
net-snmp	BZ# 1533943			
nftables	BZ# 1571968			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
nmap		BZ# 1546246 , BZ# 1573411		
nss			BZ# 1425514 , BZ# 1431210 , BZ# 1432142	
nuxwdog		BZ# 1615617		
opal-prd	BZ# 1564097			
openjpeg		BZ# 1553235		
opensc		BZ# 1547117 , BZ# 1562277 , BZ# 1562572		
openscap		BZ# 1556988		BZ# 1603347 , BZ# 1640522
openssl	BZ# 1519396			
openssl-ibmca	BZ# 1519395			
oscap-anaconda-addon				BZ# 1636847
other	BZ# 1432080 , BZ# 1609302 , BZ# 1612965 , BZ# 1627126		BZ# 1062759 , BZ# 1072107 , BZ# 1259547 , BZ# 1464377 , BZ# 1477977 , BZ# 1559615 , BZ# 1613966	BZ# 1569484 , BZ# 1571754 , BZ# 1611665 , BZ# 1633185 , BZ# 1635135
pacemaker	BZ# 1590483			
pam_pkcs11	BZ# 1578029			
pcp	BZ# 1565370			
pcs	BZ# 1427273 , BZ# 1475318	BZ# 1566382 , BZ# 1572886 , BZ# 1588667 , BZ# 1590533	BZ# 1433016	
pcsc-lite	BZ# 1516993			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
pcsc-lite-ccid	BZ# 1558258			
perl	BZ# 1557574			
perl-LDAP	BZ# 1520364			
pki-core	BZ# 1550742 , BZ# 1550786 , BZ# 1557569 , BZ# 1562423 , BZ# 1585866	BZ# 1546708 , BZ# 1549632 , BZ# 1568615 , BZ# 1580394		
powerpc-utils		BZ# 1540067 , BZ# 1592429 , BZ# 1596121 , BZ# 1628907		
procps-ng	BZ# 1518986	BZ# 1507356		
qemu-guest-agent	BZ# 1569013			
qemu-kvm			BZ# 1103193	
radvd	BZ# 1475983			
rear	BZ# 1418459 , BZ# 1496518			
resource-agents	BZ# 1470840 , BZ# 1538689 , BZ# 1568588 , BZ# 1568589		BZ# 1513957	
rhel-system-roles	BZ# 1479381		BZ# 1439896	
rpm	BZ# 1395818 , BZ# 1555326			
rsyslog	BZ# 1482819 , BZ# 1531295 , BZ# 1539193			BZ# 1553700
rt-setup	BZ# 1616038			
samba	BZ# 1558560			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
sane-backends	BZ# 1512252			
scap-security-guide	BZ# 1443551 , BZ# 1619689			BZ# 1631378
scap-workbench				BZ# 1533108
selinux-policy	BZ# 1443473 , BZ# 1460322			
sos-collector	BZ# 1481861			
sssd	BZ# 1416528		BZ# 1068725	
strongimcv			BZ# 755087	
subscription-manager	BZ# 1576423			
sudo	BZ# 1533964 , BZ# 1547974 , BZ# 1548380	BZ# 1560657		
systemd			BZ# 1284974	
systemtap	BZ# 1565773			
tss2			BZ# 1384452	
tuned	BZ# 1546598			BZ# 1576435
usbguard	BZ# 1508878		BZ# 1480100	
vdo				BZ# 1617896
vsftpd	BZ# 1479237			
wayland			BZ# 1481411	
wpa_supplicant		BZ# 1434434 , BZ# 1505404		
xorg-x11-drv-nouveau				BZ# 1624337
xorg-x11-drv-qxl				BZ# 1640918

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
xorg-x11-server	BZ# 1564632			BZ# 1624847
ypserv		BZ# 1492892		
yum	BZ# 1481220	BZ# 1528608		
yum-utils	BZ# 1497351 , BZ# 1506205			

APPENDIX C. REVISION HISTORY

Revision 0.0-8 Updated NVMe/FC-related notes. Updated Deprecated Functionality. Other various additions and updates.	Fri Nov 02 2018	Lenka Špačková
Revision 0.0-7 Release of the Red Hat Enterprise Linux 7.6 Release Notes.	Tue Oct 30 2018	Lenka Špačková
Revision 0.0-0 Release of the Red Hat Enterprise Linux 7.6 Beta Release Notes.	Wed Aug 22 2018	Lenka Špačková