



Red Hat Enterprise Linux 7

7.5 Release Notes

Release Notes for Red Hat Enterprise Linux 7.5

Red Hat Enterprise Linux 7 7.5 Release Notes

Release Notes for Red Hat Enterprise Linux 7.5

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

Table of Contents

PREFACE	15
CHAPTER 1. OVERVIEW	16
Security and Compliance	16
Performance and Efficiency	16
Platform Manageability	16
Identity Management and Access Control	16
Support for Architectures in the New Kernel Version	17
Virtualization	17
Red Hat Insights	17
Red Hat Customer Portal Labs	17
CHAPTER 2. ARCHITECTURES	19
Support for Architectures in the kernel-alt Packages	19
CHAPTER 3. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	21
KERNEL PARAMETERS	21
KERNEL PARAMETERS TO MITIGATE SPECTRE AND MELTDOWN ISSUES	22
UPDATED /PROC/SYS/NET/CORE ENTRIES	22
PART I. NEW FEATURES	24
CHAPTER 4. GENERAL UPDATES	25
In-place upgrade from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7	25
The setup package now provides a way to override unpredictable environment settings	25
CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY	26
Windows Server 2016 forest and domain functional levels now supported for trust	26
Directory Server no longer displays replication conflict entries in search results	26
OpenLDAP is now compiled with OpenSSL instead of NSS	26
Samba rebased to version 4.7.1	26
The SSSD LDAP provider can now automatically create user private groups for users	27
SSSD enrolled to an AD domain remembers the discovered AD site after the first successful connection	27
SSSD logs changes in its status to syslog	27
SSSD performance has improved	27
The pwdhash utility can now retrieve the storage scheme from the configuration directory	27
New utility to compare two Directory Server instances	27
Directory Server now supports enabling the memberOf plug-in on read-only replicas	28
Directory Server rebased to version 1.3.7.5	28
Directory Server supports additional password storage schemes	28
Directory Server now uses separate normalized DN caches for each worker thread	28
pki-core rebased to version 10.5.1	28
Certificate System supports installing CA, KRA, and OCSP subsystems with CMC	28
Certificate System supports creating instances running as a different user	29
Certificate System can now create PKCS #12 files using PBES2 with PBKDF2 key derivation	29
Certificate System CAs can now process CMC renewal requests signed by a previously issued signing certificate	29
Certificate System now uses the Mozilla NSS secure random number generator	29
Audit event changes in Certificate System	29
krb5 now includes the kdcpolicy interface	29
Certificate System now supports configurable hashing algorithms for the SKI extension	30
The pki command-line interface automatically creates a default NSS database	30
Certificate System disables weak 3DES ciphers by default	30

The Certificate System CA subsystem's OCSP provider now includes the nextUpdate field in responses	30
ding-libs rebased to version 0.6.1	30
CHAPTER 6. CLUSTERING	31
New SNMP agent to query a Pacemaker cluster	31
Support for Red Hat Enterprise Linux High Availability clusters on Amazon Web Services	31
Support for Red Hat Enterprise Linux High Availability clusters on Microsoft Azure	31
Unfencing is done in resource cleanup only if relevant parameters changed	31
The pcsd port is now configurable	31
Fencing and resource agents are now supported by AWS Python libraries and a CLI client	31
Fencing in HA setups is now supported by Azure Python libraries	31
New features added to the sbd binary.	31
sbd rebased to version 1.3.1	31
Cluster status now shows by default when a resource action is pending	32
clutter rebased to version 0.77.0	32
Support for Sybase ASE failover	32
CHAPTER 7. COMPILER AND TOOLS	33
The linuxptp package now supports active-backup bonding for clock synchronization	33
parted can now resize partitions using the resizepart command	33
binutils rebased to version 2.27	33
pcp rebased to version 3.12.2	33
Improved DWARF 5 support in various tools	34
systemtap rebased to version 3.2	34
valgrind rebased to version 3.13.0	35
ncat rebased to version 7.50	35
rsync rebased to version 3.1.2	35
tcpdump can now analyze virtio traffic	36
Vim now supports C++11 syntax highlighting	36
Vim now supports the blowfish2 encryption method	36
The IO::Socket::SSL Perl module now uses the system-wide CA certificate store by default	36
perl-DateTime-TimeZone rebased to version 1.70	36
system-config-kdump now support selecting of either automated or manual kdump memory settings when fadump is performed	37
conman rebased to version 0.2.8	37
Support for the TFTP window size option has been implemented	37
curl now supports disabling GSSAPI with SOCKS5	37
The rsync utility now copies files with their original nanosecond part of the time stamp	37
tcpdump rebased to version 4.9.2	37
OProfile support for Intel Xeon processor family extended	38
Support for Intel Xeon v4 uncore performance events in libpfm, pcp, and papi	38
Memory copying performance improved on IBM POWER architectures	38
TAI clock macro available	38
Support for selective use of 4 KiB page tables on IBM z Systems	38
More efficient glibc functions on IBM z Systems	38
The ld linker no longer incorrectly combines position-dependent and independent code	38
python-virtualenv rebased to 15.1.0	38
python-urllib3 supports IP addresses in subjectAltName	38
Support for retpolines added to GCC	38
Shenandoah garbage collector is now fully supported	39
CHAPTER 8. DESKTOP	40
GNOME Shell rebased to version 3.26	40
gnome-settings-daemon rebased to version 3.26	40

libreoffice rebased to version 5.3	40
GIMP rebased to version 2.8.22	40
Inkscape rebased to version 0.92.2	41
webkitgtk4 rebased to version 2.16	42
qt5 rebased to version 5.9.2	42
New package: qgnomeplatform	42
ModemManager rebased to version 1.6.8	43
New packages: libsmbios	43
mutter rebased to version 3.26	43
The SANE_USB_WORKAROUND environmental variable can make older scanners usable with USB3	43
The libyami package added for better video stream handling	43
netpbm rebased to version 10.79.00	44
Red Hat Enterprise Linux 7.5 supports libva	44
GStreamer now supports mp3	44
GNOME control-center rebased to version 3.26	44
New package: emacs-php-mode	44
Dutch keyboard layout provided	44
CHAPTER 9. FILE SYSTEMS	46
SMB 2 and SMB 3 now support DFS	46
File system DAX now performs better when mapping a large amount of memory	46
quotacheck is now faster on ext4	46
The CephFS kernel client is fully supported with Red Hat Ceph Storage 3	46
CHAPTER 10. HARDWARE ENABLEMENT	47
Broadcom 5880 smart card readers with the updated firmware are now supported	47
fwupd now supports Synaptics MST hubs	47
kernel-rt sources updated	47
Improved RT throttling mechanism	47
VMware Paravirtual RDMA Driver	47
opal-prd rebased to version 5.9	47
libreswan now supports NIC offloading	48
Trusted Computing Group TPM 2.0 System API library and management utilities available	48
new packages: tpm2-abrmd	48
CHAPTER 11. INSTALLATION AND BOOTING	49
Assigning mount points to existing block devices is now possible in Kickstart installations	49
The livemedia-creator utility now provides a sample Kickstart file for UEFI systems	49
New option for the network Kickstart command binding the device configuration file to the device MAC address	49
New options for Kickstart %packages allow configuring Yum timeout and number of retries	49
The Red Hat Enterprise Linux 7 ISO image can be used to create guests virtual machines on IBM z Systems	49
ARPUUPDATE option for ifcfg-* files has been introduced	50
The --noconfig option added for the rpm -V command	50
ifcfg-* files now allow you to specify a third DNS server	50
Multi-threaded xz compression in rpm-build	50
CHAPTER 12. KERNEL	51
Memory Protection Keys are now supported in later Intel processors	51
EDAC support added for Pondicherry 2 memory controllers	51
MBA is now supported	51
Swap optimizations enable fast block devices to be used as secondary memory	51
HID Wacom rebased to version 4.12	51
New livepatch functionality improves the latency and success rate of the kpatch-patch packages	51

Persistent Kernel Module Upgrade (PKMU) supported	51
The Linux kernel now supports encrypted SMB 3 connections	52
SME enabled on AMD Naples platforms	52
Support for the ie31200_edac driver	52
EDAC now supports GHES	52
CUIR enhanced scope detection is now fully supported	52
kdump allows a vmcore collection without the root file system being mounted	52
KASLR fully supported and enabled by default	52
Intel® Omni-Path Architecture (OPA) Host Software	53
noreplace-paravirt has been removed from the kernel command line parameters	53
The new EFI memmap implementation is now available on SGI UV2+ systems	53
Mounting pNFS shares with flexible file layout is now fully supported	53
CHAPTER 13. NETWORKING	54
Error handling in the output of the dhcp-script has been improved	54
Network namespace isolation has been added to ipset	54
NetworkManager now supports multiple routing tables to enable source routing	54
nftables rebased to version 0.8	54
Persistent DHCP client behavior added to NetworkManager	54
NetworkManager exposes new properties to expose team options	55
Packets mark is now reflected on replies	55
New Socket timestamping options for NTP	55
iproute2 rebased to version 4.11.0	55
The tc-pedit action now supports offset relative to Layer 2 and Layer 4	56
Features backported to iproute	56
The Geneve driver rebased to version 4.12	56
A control switch added for VXLAN and GENEVE offloading	56
unbound rebased to version 1.6.6	56
DHCP now supports standard dynamic DNS updates	57
DDNS now supports additional algorithms	57
IPTABLES_SYSCTL_LOAD_LIST now supports the sysctl.d files	58
SCTP now supports MSG_MORE	58
MACsec rebased to version 4.13	58
Enhanced performance when using the mlx5 driver in Open vSwitch	58
The Netronome NFP Ethernet driver now supports the representor netdev feature	58
Support for offloading TC-Flower actions	58
DNS stub resolver improvements	59
CHAPTER 14. SECURITY	60
LUKS-encrypted removable storage devices can be now automatically unlocked using NBDE	60
new package: clevis-systemd	60
OpenSCAP can be now integrated into Ansible workflows	60
SECCOMP_FILTER_FLAG_TSYNC enables synchronization of calling process threads	60
nss rebased to version 3.34	60
SSLv3 disabled in mod_ssl	60
Libreswan now supports split-DNS configuration for IKEv2	61
libreswan now supports AES-GMAC for ESP	61
openssl-ibmca rebased to 1.4.0	61
opencryptoki rebased to 3.7.0	61
atomic scan with configuration_compliance enables creating security-compliant container images at build time	61
tang-nagios enables Nagios to monitor Tang	61
clevis now logs privileged operations	62

PK11_CreateManagedGenericObject() has been added to NSS to prevent memory leaks in applications	62
OpenSSH now supports openssl-ibmca and openssl-ibmpkcs11 HSMs	62
cgroup_seclabel enables fine-grained access control on cgroups	62
The boot process can now unlock encrypted devices connected by network	62
SELinux now supports InfiniBand object labeling	62
libica rebased to 3.2.0	63
SELinux now supports systemd No New Privileges	63
Libreswan rebased to version 3.23	63
libreswan now supports IKEv2 MOBIKE	64
scap-workbench rebased to version 1.1.6	64
OpenSCAP is now able to generate results for DISA STIG Viewer	64
selinux-policy no longer contains permissive domains	64
audit rebased to version 2.8.1	64
OpenSC now supports the SCE7.0 144KDI CAC Alt. tokens	65
CHAPTER 15. SERVERS AND SERVICES	66
Leftover dbus processes	66
dbus rebased to version 1.10	66
tuned rebased to version 2.9.0	66
chrony rebased to version 3.2	66
SNMP page counting can be now disabled in CUPS	66
CUPS can be set to use only ciphers from TLS version 1.2 or later	67
The squid packages now provide the kerberos_ldap_group helper	67
OpenIPMI rebased to version 2.0.23	67
Overview of changes from freeIPMI 1.2.9 to freeIPMI 1.5.7	67
A new clear_env option available in PHP FPM pool configuration	67
CHAPTER 16. STORAGE	68
Data Deduplication and Compression with VDO	68
New boom utility for managing LVM snapshot and image boot entries	68
DM Multipath no longer requires reservation keys in advance	68
New property parameter supported in blacklist and blacklist_exception sections of multipath.conf	68
smartmontools now support NVMe devices	69
Support for DIF/DIX (T10 PI) on specified hardware	69
File system Direct Access (DAX) and device DAX now support huge pages	70
fsadm can now grow and shrink LUKS-encrypted LVM volumes	70
CHAPTER 17. SYSTEM AND SUBSCRIPTION MANAGEMENT	71
cockpit rebased to version 154	71
Users of yum-utils now can perform actions prior to transactions	71
yum can disable creation of per-user cache as a non-root user	71
yum-builddep now allows to define RPM macros	71
subscription-manager now displays the host name upon registration	71
A subscription-manager plugin now runs with yum-config-manager	72
subscription-manager now protects all product certificates in /etc/pki/product-default/	72
rhnmigrateclassic-to-rhsm now automatically enables the subscription-manager and product-id yum plugins	72
subscription-manager now automatically enables the subscription-manager and product-id yum plugins	72
subscription-manager-cockpit replaces subscription functionality in cockpit-system	72
virt-who logs where the host-guest mapping is sent	72
virt-who now provides configuration error information	72
reposync now by default skips packages whose location falls outside the destination directory	73
CHAPTER 18. VIRTUALIZATION	74
KVM virtualization on IBM z Systems	74

KVM virtualization supported on IBM POWER9	74
KVM virtualization supported on IBM POWER8	74
NVIDIA GPU devices can now be used by multiple guests simultaneously	74
KASLR for KVM guests	74
Parallel decompression of OVA files supported	75
SMAP now supported on Cannonlake guests	75
libvirt rebased to 3.9.0	75
virt-manager rebased to 1.4.3	75
virt-what rebased to version 1.18	76
tboot rebased to version 1.96	76
virt-v2v can convert VMware guests with snapshots	76
virt-rescue enhanced	77
virt-v2v now converts Linux guests encrypted with LUKS	77
CAT support added to libvirt on specific CPU models	77
PTP device added to improve time synchronization of KVM guests	77
CHAPTER 19. RED HAT ENTERPRISE LINUX 7.5 FOR ARM	78
19.1. NEW FEATURES AND UPDATES	78
19.2. KERNEL CONFIGURATION CHANGES	78
HARDWARE ENABLEMENT	78
CORE KERNEL SUPPORT	80
19.3. SUPPORT IN RED HAT SATELLITE	82
19.4. KNOWN ISSUES	82
19.5. BUG FIXES	83
CHAPTER 20. RED HAT ENTERPRISE LINUX 7.5 FOR IBM POWER LE (POWER9)	85
20.1. NEW FEATURES AND UPDATES	85
20.2. KERNEL CONFIGURATION CHANGES	86
HARDWARE ENABLEMENT	86
CORE KERNEL SUPPORT	87
20.3. SUPPORT IN RED HAT SATELLITE	89
20.4. KNOWN ISSUES	89
20.5. BUG FIXES	91
CHAPTER 21. ATOMIC HOST AND CONTAINERS	92
Red Hat Enterprise Linux Atomic Host	92
CHAPTER 22. RED HAT SOFTWARE COLLECTIONS	93
PART II. NOTABLE BUG FIXES	94
CHAPTER 23. GENERAL UPDATES	95
runc notifies systemd about user-specified CPU quota limits	95
Segmentation faults in applications because of only non-existent paths in LD_LIBRARY_PATH no longer happen	95
The setup package now creates the tape group with the correct group number	95
CHAPTER 24. AUTHENTICATION AND INTEROPERABILITY	96
The IdM LDAP server no longer becomes unresponsive when resolving an AD user takes a long time	96
Application configuration snippets in /etc/krb5.conf.d/ are now automatically read in existing configurations	96
pam_mkhomedir can now create home directories under /	96
Kerberos operations depending on KVNO in the keytab file no longer fail when a RODC is used	96
krb5 properly displays errors about PKINIT misconfiguration in single-realm KDC environments	96
Certificate System no longer incorrectly logs ROLE_ASSUME audit events	97
Updated attributes in CERT_STATUS_CHANGE_REQUEST_PROCESSED audit log event	97

Signed audit log verification now works correctly	97
Certificate System now validates the banner file	97
The TPS subsystem no longer fails when performing a symmetric key changeover on a HSM	97
Certificate System CAs no longer display an error when handing subject DN's without a CN component	97
The pki-server-upgrade utility no longer fails if target files are missing	98
The Certificate System CA key replication now works correctly	98
Certificate System no longer fails to import PKCS #12 files	98
The TPS user interface now displays the token type and origin fields	98
Certificate System issued certificates with an expiration date later than the expiration date of the CA certificate	98
CA certificates without SKI extension no longer causes issuance failures	98
Certificate System correctly logs the user name in CMC request audit events	98
The Directory Server trivial word check password policy now works as expected	98
The pkidestroy utility now fully removes instances that are started by the pki-tomcatd-nuxwdog service	99
The Certificate System deployment archive file no longer contains passwords in plain text	99
ACIs with the targetfilter keyword work correctly	99
Directory Server searches with a scope set to one have been fixed	99
Clear error message when sending TLS data to a non-LDAPS port	99
Directory Server no longer logs an error if not running the cleanallruv task	99
Using a large number of CoS templates no longer slow down the virtual attribute processing time	100
Directory Server now handles binds during an online initialization correctly	100
The dirsrv@.service meta target is now linked to multi-user.target	100
The memberOf plug-in now logs all update attempts of the memberOf attribute	100
The Directory Server password policies now work correctly	100
A buffer overflow has been fixed in Directory Server	100
Directory Server now sends the password expired control during grace logins	100
An unnecessary global lock has been removed from Directory Server	100
Replication now works correctly with TLS client authentication and FIPS mode enabled	101
Directory Server now correctly sets whether virtual attributes are operational	101
Backup now succeeds if replication was enabled and a changelog file existed	101
Certificate System updates the revocation reason correctly	101
A race condition has been fixed in the Certificate System clone installation process	101
Certificate System now uses strong ciphers by default	101
The pkispawn utility no longer displays incorrect errors	102
The Certificate System profile configuration update method now correctly handles backslashes	102
CHAPTER 25. CLUSTERING	103
Pacemaker correctly implements fencing and unfencing for Pacemaker remote nodes	103
Pacemaker now probes guest nodes	103
The pcs resource cleanup command no longer generates unnecessary cluster load	103
Warning generated when user specifies action attribute for stonith device	103
It is now possible to enable stonith agent debugging without specifying the --force flag	103
The fence_ilo3 resource agent no longer has a default value of cycle for the action parameter	103
Pacemaker no longer starts up when sbd is enabled but not started successfully by systemd	104
A fenced node in an 'sbd' setup now shuts down reliably	104
IPaddr2 resource agent now finds NIC for IPv6 addresses with 128 netmask	104
portblock agent no longer yields excessive unnecessary messages	104
/var/run/resource-agents directory now persists across reboots	104
CHAPTER 26. COMPILER AND TOOLS	105
Package selection now works in system-config-kickstart	105
NVMe devices no longer show up as Unknown in parted and Anaconda	105
DBD::MySQL now sends and receives smaller integers correctly on big-endian platforms	105

The version Perl module now supports tainted input and tainted version objects	105
The HTTP::Daemon Perl module now supports IPv6	105
GDB shows inline function names in breakpoint listing	105
Relocation failures at module load time due to wrong GCC alignment fixed	105
The istream::sentry object from the gcc C++ standard library no longer throws exceptions	105
Multiple fixes in gdb on IBM Power	106
GDB no longer crashes when dumping core from a process that terminates	106
GDB can again dump memory protected by the VM_DONTDUMP flag	106
Programs using the CLONE_PTRACE flag on threads now run under strace	106
exiv2 rebased to version 0.26	106
gssproxy fixed to properly update ccaches	107
gcc on the little-endian variant of IBM Power Systems architecture no longer creates unused stack frames	107
Several bugs fixed in gssproxy	107
The BFD library regains the ability to convert binary addresses to source code positions	107
Applications using vector registers for passing arguments work again	107
curl now properly resets the HTTP authentication state	107
The strip utility works again	107
Importing python modules generated by f2py now works properly	107
mailx is not encoding multi-byte subjects properly	108
The --all-logs option now works as expected in sosreport	108
Python scripts can now correctly connect to HTTPS servers through a proxy, while explicitly setting the port	108
CHAPTER 27. DESKTOP	109
Stylus of Dell Canvas 27 fixed	109
llvmpipe crashes on IBM Power Systems	109
CHAPTER 28. FILE SYSTEMS	110
NFS shares no longer become unresponsive after a TCP connection is closed	110
CHAPTER 29. HARDWARE ENABLEMENT	111
genwqe-tools updated for IBM Power Systems ppc64 and ppc64le architectures	111
Hardware utility tools now correctly identify recently released hardware	111
CHAPTER 30. INSTALLATION AND BOOTING	112
The installer no longer crashes when you select an incomplete IMSM RAID array during manual partitioning	112
Installer now accepts additional time zone definitions in Kickstart files	112
Proxy configuration set up using a boot option now works correctly in Anaconda	112
FIPS mode now supports loading files over HTTPS during installation	112
Network scripts now correctly update /etc/resolv.conf	112
Files with the .old extension are now ignored by network scripts	113
Bridge devices no longer fail to obtain an IP address	113
The rhel-dmesg service can now be disabled correctly	113
CHAPTER 31. KERNEL	114
kdump can now capture a vmcore with nokaslr set	114
MPOL_PREFERRED policy now works with Transparent Huge Pages (THP) with optimal performance	114
A cgroups deadlock has been fixed	114
System no longer becomes unresponsive when DM thin provisioning is used on top of a loop device	114
KASLR now no longer causes mirroring of kernel memory to non-mirrored regions	114
Users now receive message with prompt to remove white space characters in the /etc/kdump.conf	114
An application with large .bss segment on IBM POWER Systems will no longer cause random segmentation faults	114
Kernel no longer consumes excessive amounts of resources to calculate load	114
Cpuset is now able to restore the effective CPU mask after a pair of offline and online events	115

Access to /proc/[pid]/maps is now significantly faster	115
fadump no longer fails to restart	115
makedumpfile can now map page table entries correctly	115
Asymmetric groups are used for overlapping scheduling domains	115
The KASLR no longer causes kernel to become unresponsive while booting the system	115
Unplugging a Wacom tablet with ExpressKeys no longer causes the operating system to reboot	116
Setting memory.kmem.limit_in_bytes no longer causes a problem when removing that memory cgroup later	116
The sha1-avx2 encryption algorithm is now re-enabled	116
VXLAN rebased to version 4.14	116
CHAPTER 32. NETWORKING	117
Network operation persists when ip6mr unregisters an already unregistered device	117
Sending big files through VTI no longer fails	117
L2TP with IPv6 encapsulation now works in name space	117
Flushing ARP entries no longer fails	117
Using cls_matchall with classful queue disciplines no longer causes the kernel to crash	117
ICMP error packets are no longer lost when a user connects to a closed SCTP port	117
SCTP now selects the right source address	117
Device reference held by iptables CLUSTERIP target is now properly released on namespace deletion	118
The nftables configuration files are no longer publicly readable	118
The Ready to read events are now correctly sent to an application when SENDER_DRY_EVENTS is enabled	118
SCTP statistics now available	118
The firewalld service daemon no longer hangs in the rmmod process	118
CHAPTER 33. SECURITY	119
When firewalld starts, net.netfilter.nf_conntrack_max is no longer reset to default if its configuration exists	119
Tomcat can now be started using tomcat-jsvc with SELinux in enforcing mode	119
SELinux now allows vdsmd to communicate with lldpad	119
OpenSSH servers without Privilege Separation no longer crash	119
The clevis luks bind command no longer fails with the DISA STIG-compliant password policy	119
WinSCP 5.10 now works properly with OpenSSH	119
SFTP no longer allows to create zero-length files in read-only mode	119
CHAPTER 34. SERVERS AND SERVICES	120
Internal buffer locks no longer cause deadlocks in libdb	120
Weekly log rotations are now triggered more predictably	120
ghostscript no longer crashes while processing large PDF files	120
Converting large PDF files to PNG with ghostscript no longer fails	120
krfb no longer crashes when unable to bind to an IPv6 port	120
mod_nss properly detects the threading model in Apache to improve performance	120
atd no longer runs with 100% CPU utilization nor fills system log	120
ReaR now provides a more helpful error message when grub2-efi-x64-modules is missing	120
ReaR no longer fails to determine disk size during a mkrescue operation	121
ReaR no longer requires dosfsck and efibootmgr on non-UEFI systems	121
ReaR no longer fails with NetBackup and has more reliable network configuration	121
ReaR recovery no longer fails when backup integrity checking is enabled	121
CHAPTER 35. STORAGE	122
DM Multipath no longer crashes when adding a feature to an empty string	122
I/O operations no longer hang with RAID1	122
CHAPTER 36. SYSTEM AND SUBSCRIPTION MANAGEMENT	123
Yum no longer crashes in certain nss and nspr update scenario	123

The fastestmirror plug-in now orders mirrors before the metadata download	123
The package-cleanup script no longer removes package dependencies of non-duplicates	123
rhnsd.pid is now writable only by the owner	123
rhnsd now correctly reports system reboots to Satellite	123
The rpm rhnlib -qi command now refers to the current upstream project website	123
Kernel installations using rhnsd complete successfully	123
rhnsd no longer modifies permissions on files in /var/cache/yum/	123
subscription-manager reports an RPM package if its vendor contains non-UTF8 characters	124
subscription-manager now works with proxies that expect the Host header	124
subscription-manager assigns valid IPv4 addresses to network.ipv4_address even if initial DNS resolution fails	124
virt-who ensures that provided options fit the same virtualization type	124
virt-who configuration no longer resets on upgrade or reinstall	124
virt-who now reads the 'address' field provided by RHEVM to discover and report the correct host name	124
CHAPTER 37. VIRTUALIZATION	125
Guests no longer shut down unexpectedly during reboot	125
Guests accessed using a serial console no longer become unresponsive	125
virt-v2v now warns about not converting PCI passthrough devices	125
When importing OVAs, virt-v2v now parses MAC addresses	125
PART III. TECHNOLOGY PREVIEWS	126
CHAPTER 38. GENERAL UPDATES	127
The systemd-importd VM and container image import and export service	127
CHAPTER 39. AUTHENTICATION AND INTEROPERABILITY	128
Use of AD and LDAP sudo providers	128
DNSSEC available as Technology Preview in IdM	128
Identity Management JSON-RPC API available as Technology Preview	128
The Custodia secrets service provider is now available	128
Containerized Identity Management server available as Technology Preview	129
CHAPTER 40. CLUSTERING	130
The pcs tool now manages bundle resources in Pacemaker	130
New fence-agents-heuristics-ping fence agent	130
Heuristics supported in corosync-qdevice as a Technology Preview	130
CHAPTER 41. DESKTOP	131
Wayland available as a Technology Preview	131
Fractional Scaling available as a Technology Preview	131
CHAPTER 42. FILE SYSTEMS	132
ext4 and XFS file systems now support DAX	132
pNFS block layout is now available	132
pNFS SCSI layout is now available for client and server	132
OverlayFS	132
Btrfs file system	133
New package: ima-evm-utils	133
CHAPTER 43. HARDWARE ENABLEMENT	134
LSI Syncro CS HA-DAS adapters	134
tss2 enables TPM 2.0 for IBM Power LE	134
ibmvnic Device Driver	134
CHAPTER 44. KERNEL	135

Heterogeneous memory management included as a Technology Preview	135
criu rebased to version 3.5	135
kexec as a Technology Preview	135
kexec fast reboot as a Technology Preview	135
Unprivileged access to name spaces can be enabled as a Technology Preview	135
SCSI-MQ as a Technology Preview in the qla2xxx driver	135
NVMe over Fibre Channel is now available as a Technology Preview	136
perf cqm has been replaced by resctrl	136
CHAPTER 45. REAL-TIME KERNEL	138
The SCHED_DEADLINE scheduler class as Technology Preview	138
CHAPTER 46. NETWORKING	139
Cisco usNIC driver	139
Cisco VIC kernel driver	139
Trusted Network Connect	139
SR-IOV functionality in the qlcnict driver	139
The libnftnl and nftables packages	139
The flower classifier with off-loading support	139
CHAPTER 47. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE	140
Red Hat Enterprise Linux System Roles	140
CHAPTER 48. SECURITY	141
USBGuard enables blocking USB devices while the screen is locked as a Technology Preview	141
pk12util can now import certificates signed with RSA-PSS	141
Support for certificates signed with RSA-PSS in certutil has been improved	141
NSS is now able to verify RSA-PSS signatures on certificates	141
SECCOMP can be now enabled in libreswan	141
CHAPTER 49. STORAGE	143
Multi-queue I/O scheduling for SCSI	143
Targetd plug-in from the libStorageMgmt API	143
Support for Data Integrity Field/Data Integrity Extension (DIF/DIX)	143
CHAPTER 50. VIRTUALIZATION	144
USB 3.0 support for KVM guests	144
Select Intel network adapters now support SR-IOV as a guest on Hyper-V	144
No-IOMMU mode for VFIO drivers	144
virt-v2v can now use vmx configuration files to convert VMware guests	144
virt-v2v can convert Debian and Ubuntu guests	144
Virtio devices can now use vIOMMU	144
virt-v2v converts VMWare guests faster and more reliably	144
Open Virtual Machine Firmware	145
PART IV. DEVICE DRIVERS	146
CHAPTER 51. NEW DRIVERS	147
Storage Drivers	147
Network Drivers	147
Graphics Drivers and Miscellaneous Drivers	147
CHAPTER 52. UPDATED DRIVERS	148
Storage Driver Updates	148
Network Driver Updates	148
Graphics Driver and Miscellaneous Driver Updates	149

CHAPTER 53. DEPRECATED FUNCTIONALITY	150
Python 2 has been deprecated	150
LVM libraries and LVM Python bindings have been deprecated	150
Mirrored mirror log has been deprecated in LVM	150
Deprecated packages related to Identity Management and security	150
Support for earlier IdM servers and for IdM replicas at domain level 0 will be limited	152
Bug-fix only support for the nss-pam-ldapd and NIS packages in the next major release of Red Hat Enterprise Linux	152
Use the Go Toolset instead of golang	152
mesa-private-llvm will be replaced with llvm-private	153
libdbi and libdbi-drivers have been deprecated	153
Ansible deprecated in the Extras channel	153
signtool has been deprecated	153
TLS compression support has been removed from nss	154
Public web CAs are no longer trusted for code signing by default	154
Sendmail has been deprecated	154
dmraid has been deprecated	154
Automatic loading of DCCP modules through socket layer is now disabled by default	154
rsyslog-libdbi has been deprecated	154
The inputname option of the rsyslog imudp module has been deprecated	154
SMBv1 is no longer installed with Microsoft Windows 10 and 2016 (updates 1709 and later)	154
FedFS has been deprecated	155
Btrfs has been deprecated	155
tcp_wrappers deprecated	155
nautilus-open-terminal replaced with gnome-terminal-nautilus	155
sslwrap() removed from Python	155
Symbols from libraries linked as dependencies no longer resolved by ld	155
Windows guest virtual machine support limited	156
libnetlink is deprecated	156
S3 and S4 power management states for KVM have been deprecated	156
The Certificate Server plug-in udnPwDirAuth is discontinued	156
Red Hat Access plug-in for IdM is discontinued	156
The Ipsilon identity provider service for federated single sign-on	156
Several rsyslog options deprecated	156
Deprecated symbols from the memkind library	156
Options of Sockets API Extensions for SCTP (RFC 6458) deprecated	157
Managing NetApp ONTAP using SSLv2 and SSLv3 is no longer supported by libstorageMgmt	157
dconf-dbus-1 has been deprecated and dconf-editor is now delivered separately	157
FreeRADIUS no longer accepts Auth-Type := System	157
Deprecated Device Drivers	158
Deprecated Adapters	160
The libcxgb3 library and the cxgb3 firmware package have been deprecated	165
SFN4XXX adapters have been deprecated	165
Software-initiated-only FCoE storage technologies have been deprecated	165
Containers using the libvirt-lxc tooling have been deprecated	166
PART V. KNOWN ISSUES	167
CHAPTER 54. AUTHENTICATION AND INTEROPERABILITY	168
A crash is reported after an unsuccessful lightweight CA key retrieval	168
OpenLDAP causes programs to fail immediately in case of incorrect configuration	168
OpenLDAP reports failures when CACertFile or CACertDir point to an invalid location	168
OpenLDAP does not update TLS configuration after inconsistent changes in cn=config	168

Identity Management terminates connections unexpectedly	168
Directory Server can terminate unexpectedly during shutdown	168
CHAPTER 55. CLUSTERING	170
Data corruption occurs on RAID 10 reshape on top of VDO with el7 kernel.	170
CHAPTER 56. COMPILER AND TOOLS	171
Memory consumption of applications using libcurl grows with each TLS connection	171
OProfile and perf can not sample events on 2nd generation Intel Xeon Phi processors when NMI watchdog is disabled	171
CHAPTER 57. DESKTOP	172
Cannot install downloaded RPM files from Nautilus	172
Caps Lock LED status	172
Inconsistent GNOME Shell versions	172
Uninstall the 32-bit version of flatpak	172
GNOME downgrade does not work	172
Wayland ignores keyboard grabs issued by X11 applications, such as virtual machines viewers	172
Superuser should not run graphical sessions	173
Keyboard not working in VM browsed by remote-viewer and virt-viewer	173
gnome-system-log does not work on Wayland	173
GUI screen is shown incorrectly	173
xrandr fails to provide some video modes	173
radeon fails to reset hardware correctly	173
nouveau fails to load Nvidia secboot firmware	174
Xchat status icon disappears from Top Icons panel	174
GDM does not activate hotplugged monitors	174
Wacom Expresskeys Remote not detected as tablet	174
Synaptics dependency removes xorg-x11-drivers	174
T470s docking station jack does not work on resume	174
Screen occasionally turns off when xrandr is executed	175
HDMI and DP for 8th generation Intel Core processors not enumerating sound inputs	175
Tray icons are non-responsive for auto-started applications	175
Inconsistent panel color on login screen	175
Additional displays are mirrored after attaching a VM guest	175
CHAPTER 58. INSTALLATION AND BOOTING	176
Selecting the Lithuanian language causes the installer to crash	176
oscap-anaconda-addon fails to remediate when installing in TUI using Kickstart	176
The grub2-mkimage command fails on UEFI systems by default	176
Kernel panic during RHEL 7.5 installation on HPE BL920s Gen9 systems	176
The READONLY=yes option is not sufficient to configure a read-only system	176
CHAPTER 59. KERNEL	178
Security patches addressing Spectre and Meltdown issues can cause performance loss	178
The KSC does not support the xz compression	178
The update of megaraid_sas can lead to a performance decrease	178
qedi fails to bind to the iSCSI PCIe function if qede is loaded	178
radeon causes a kernel panic	179
Kdump kernel fails to boot after a CPU hot add or hot remove operation	179
CHAPTER 60. NETWORKING	180
Verification of signatures using the MD5 hash algorithm is disabled in Red Hat Enterprise Linux 7	180
CHAPTER 61. SECURITY	181

NSS accept malformed RSA PKCS#1 v1.5 signatures made with an RSA-PSS key	181
Authentication using ssh-agent not from OpenSSH fails	181
Parsing of OpenSSH public keys is more strict	181
SCAP Workbench fails to generate results-based remediations from tailored profiles	181
Clevis can log spurious Device is not initialized error messages	181
Libreswan is not working properly with seccomp=enabled on all configurations	181
OpenSCAP RPM verification rules do not work correctly with VM and container file systems	181
Firefox and other applications using NSS become unresponsive when a smart card is inserted	182
CHAPTER 62. SERVERS AND SERVICES	183
No clear indication of profile activation error in the Tuned service	183
db_hotbackup -c should be used with caution	183
Setting ListenStream= options in rpcbind.socket causes systemd-logind to fail and SSH connections to be delayed	183
ReaR recovery process fails on non-UEFI systems with the grub2-efi-x64 package installed	183
ISO images generated by ReaR with Linux TSM fail to work	183
Unexpected problems with the dbus rebase	183
CHAPTER 63. STORAGE	184
The kexec -e command might cause storage errors with advanced storage controllers	184
LVM does not support event-based autoactivation of incomplete volume groups	184
CHAPTER 64. VIRTUALIZATION	185
Guests reporting cmt, mbmt, or mbml perf events fail to boot	185
APPENDIX A. COMPONENT VERSIONS	186
APPENDIX B. LIST OF BUGZILLAS BY COMPONENT	187
APPENDIX C. REVISION HISTORY	199

PREFACE

Red Hat Enterprise Linux (RHEL) minor releases are an aggregation of individual security, enhancement, and bug fix errata. The *Red Hat Enterprise Linux 7.5 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 and Compliance

- Security improvements and usability enhancements for cloud and remotely hosted systems that can more securely unlock Network Bound Disk Encrypted devices at boot-time. This eliminates the need for manual intervention during the often inconveniently-timed boot process.
- The integration of Red Hat Ansible Automation with OpenSCAP, which enhances the ease of automating the remediation of compliance issues and enables administrators to more efficiently deploy policies across their environment.
- Compliance improvements for accurate timestamping and synchronization needs with the addition of failover with bonding interfaces for Precision Time Protocol (PTP) and Network Time Protocol (NTP).

See [Chapter 14, Security](#) and [Chapter 7, Compiler and Tools](#) for more information.

Performance and Efficiency

- The introduction of Virtual Data Optimizer (VDO), designed to reduce data redundancy through inline deduplication and compression of primary storage. The incorporated data reduction technology helps to increase storage efficiency and reduce the cost of storage.
- Distributed File System (DFS) supported in Server Message Block (SMB) protocol versions 2 and 3. This enables a Windows system administrator to combine multiple SMB file systems into a single virtual file system.

For details, see [Chapter 16, Storage](#) and [Chapter 9, File Systems](#).

Platform Manageability

- Enhanced usability of the **Cockpit** administrator console, which is designed to simplify the interface for managing storage, networking, containers, services, and more for individual systems.
- A new utility, **boom**, which provides a command-line tool and an API for improved management of boot loader entries for LVM snapshots and images.

For details, see [Chapter 17, System and Subscription Management](#) and [Chapter 16, Storage](#).

Identity Management and Access Control

- Windows Server 2016 forest and domain functional levels are now supported for a cross-forest trust with Identity Management.
- The handling of replication conflict entries in Directory Server has been enhanced.
- The OpenLDAP suite is now compiled with the OpenSSL library instead of the Mozilla implementation of Network Security Services (Mozilla NSS).
- The samba packages have been upgraded to upstream version 4.7.1. Notably, the Samba suite in Red Hat Enterprise Linux is now using the SMB protocol version 3 by default.
- Multiple enhancements for the System Security Services Daemon (SSSD) have been introduced.

- The performance and stability of the Active Directory integration solutions provided by Identity Management have been enhanced.

For details, see [Chapter 5, Authentication and Interoperability](#).

Support for Architectures in the New Kernel Version

Red Hat Enterprise Linux 7.5 is distributed with the kernel-alt packages, which include kernel version 4.14. This kernel version provides support for the following architectures:

- 64-bit ARM
- IBM POWER9 (little endian)
- IBM z Systems

For details, see [Chapter 2, Architectures](#).

Virtualization

- KVM virtualization is now supported on IBM POWER8 systems. In addition, this update introduces support for KVM virtualization on the IBM POWER9 (little-endian) and IBM z Systems architectures. However, these require the use of kernel version 4.14, provided by the kernel-alt packages.

For details, see [Chapter 18, Virtualization](#) and [Chapter 2, Architectures](#).

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:

- [Spectre And Meltdown Detector](#)
- [Registration Assistant](#)
- [Red Hat Code Browser](#)
- [Kickstart Generator](#)
- [Red Hat Product Certificates](#)
- [Load Balancer Configuration Tool](#)
- [Red Hat Network \(RHN\) System List Exporter](#)
- [Log Reaper](#)

- [Product Life Cycle Checker](#)
- [JVM Options Configuration Tool](#)

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 7.5 is available as a single kit on the following architectures: ^[1]

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

Support for Architectures in the kernel-alt Packages

Red Hat Enterprise Linux 7.5 is distributed with the kernel-alt packages, which include kernel version 4.14. This kernel version provides support for the following architectures:

- 64-bit ARM
- IBM POWER9 (little endian) ^[5]
- IBM z Systems

The following table provides an overview of architectures supported by the two kernel versions available in Red Hat Enterprise Linux 7.5:

Table 2.1. Architectures Supported in Red Hat Enterprise Linux 7.5

Architecture	Kernel version 3.10	Kernel version 4.14
64-bit AMD and Intel	yes	no
64-bit ARM	no	yes
IBM POWER7 (big endian)	yes	no
IBM POWER8 (big endian)	yes	no
IBM POWER8 (little endian)	yes	no
IBM POWER9 (little endian)	no	yes
IBM z System	yes[a]	yes (Structure A)
[a] The 3.10 kernel version does not support KVM virtualization and containers on IBM z Systems. Both of these features are supported on the 4.14 kernel on IBM z Systems - this offering is also referred to as Structure A.		

For more information, see [Chapter 19, Red Hat Enterprise Linux 7.5 for ARM](#) and [Chapter 20, Red Hat Enterprise Linux 7.5 for IBM Power LE \(POWER9\)](#).

[1] Note that the Red Hat Enterprise Linux 7.5 installation is supported only on 64-bit hardware. Red Hat Enterprise Linux 7.5 is able to run 32-bit operating systems, including previous versions of Red Hat Enterprise Linux, as virtual machines.

[2] Red Hat Enterprise Linux 7.5 POWER8 (big endian) are currently supported as KVM guests on Red Hat Enterprise Linux 7.5 POWER8 systems that run the KVM hypervisor.

[3] Red Hat Enterprise Linux 7.5 POWER8 (little endian) is currently supported as a KVM guest on Red Hat Enterprise Linux 7.5 POWER8 systems that run the KVM hypervisor. In addition, Red Hat Enterprise Linux 7.5 POWER8 (little endian) guests are supported on Red Hat Enterprise Linux 7.5 POWER9 systems that run the KVM hypervisor in POWER8-compatibility mode on version 4.14 kernel using the kernel-alt package.

[4] Red Hat Enterprise Linux 7.5 for z Systems (both the 3.10 kernel version and the 4.14 kernel version) is currently supported as a KVM guest on Red Hat Enterprise Linux 7.5 for z Systems hosts that run the KVM on version 4.14 kernel using the kernel-alt package.

[5] Red Hat Enterprise Linux 7.5 POWER9 (little endian) is currently supported as a KVM guest on Red Hat Enterprise Linux 7.5 POWER9 systems that run the KVM hypervisor 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 7.5. 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

amd_iommu_intr = [HW,X86-64]

Specifies one of the following **AMD IOMMU** interrupt remapping modes.

legacy - Use legacy interrupt remapping mode.

vapic - Use virtual APIC mode, which allows **IOMMU** to inject interrupts directly into guest. This mode requires **kvm-amd.avic=1**, which is default when **IOMMU HW** support is present.

debug_pagealloc = [KNL]

When **CONFIG_DEBUG_PAGEALLOC** is set, this parameter enables the feature at boot time. It is disabled by default. To avoid allocating huge chunk of memory for **debug pagealloc** do not enable it at boot time, and the operating system will work similarly as with the kernel built without **CONFIG_DEBUG_PAGEALLOC**.

Use **debug_pagealloc = on** to enable the feature.

fttrace_graph_max_depth = uint[FTRACE]

This parameter is used with the function graph tracer. It defines the maximum depth it will trace into a function. Its value can be changed at run time by the **max_graph_depth file** in the **tracefs** tracing directory.

The default values is 0, which means that no limit is set.

init_pkru = [x86]

Specifies the default memory protection keys rights register contents for all processes.

The default value is 0x55555554, which disallows access to all but pkey 0. You can override the value in the **debugfs** file system after boot.

nopku = [x86]

Disables the Memory Protection Keys CPU feature found in some Intel CPUs.

mem_encrypt = [X86-64]

Provides AMD Secure Memory Encryption (SME) control. The valid arguments are: on, off.

The default setting depends on kernel configuration option:

on : **CONFIG_AMD_MEM_ENCRYPT_ACTIVE_BY_DEFAULT=y**

off : **CONFIG_AMD_MEM_ENCRYPT_ACTIVE_BY_DEFAULT=n**

mem_encrypt=on: Activate SME

mem_encrypt=off: Do not activate SME

KERNEL PARAMETERS TO MITIGATE SPECTRE AND MELTDOWN ISSUES

kpti = [X86-64]

Enables kernel page table isolation.

nopti = [X86-64]

Disables kernel page table isolation.

nospectre_v2 = [X86]

Disables all mitigations for the Spectre variant 2 (indirect branch speculation) vulnerability. The operating system may allow data leaks with this option, which is equivalent to spectre_v2=off.

spectre_v2 = [X86]

Controls mitigation of Spectre variant 2 (indirect branch speculation) vulnerability.

The valid arguments are: on, off, auto.

on: unconditionally enable

off: unconditionally disable

auto: kernel detects whether your CPU model is vulnerable

Selecting **on** will, and **auto** may, choose a mitigation method at run time according to the CPU, the available microcode, the setting of the CONFIG_RETPOLINE configuration option, and the compiler with which the kernel was built.

You can also select specific mitigations manually:

retpoline: replaces indirect branches

ibrs: Intel: Indirect Branch Restricted Speculation (kernel)

ibrs_always: Intel: Indirect Branch Restricted Speculation (kernel and user space)

Not specifying this option is equivalent to spectre_v2=auto.

UPDATED /PROC/SYS/NET/CORE ENTRIES

dev_weight_rx_bias

The **RPS** processing, for example **RFS** and **aRFS**, is competing with the registered **NAPI** poll function of the driver for the per softirq cycle **netdev_budget**.

This parameter influences the proportion of the configured **netdev_budget** that is spent on **RPS** based packet processing during RX softirq cycles. It also makes current **dev_weight** adaptable for asymmetric CPU needs on receiving on transmitting side of the network stack.

This parameter is effective on a per CPU basis. Determination is based on **dev_weight**, and it is calculated in multiplicative way (**dev_weight** * **dev_weight_rx_bias**). The default value is 1.

dev_weight_tx_bias

This parameter scales the maximum number of packets that can be processed during a TX softirq cycle.

It is effective on a per CPU basis, and allows scaling of current **dev_weight** for asymmetric net stack processing needs. Make sure to avoid making TX softirq processing a CPU hog.

Determination is based on **dev_weight**, and it is calculated in multiplicative way ($\text{dev_weight} * \text{dev_weight_rx_bias}$). The default value is 1.

PART I. NEW FEATURES

This part documents new features and major enhancements introduced in Red Hat Enterprise Linux 7.5.

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)

The setup package now provides a way to override unpredictable environment settings

The setup package now provides and sources the **sh.local** and **csh.local** files for overrides of environment variables from the **/etc/profile.d** directory, which is sourced last. Previously, an undefined order could result in unpredictable environment settings, especially when multiple scripts changed the same environment variable. (BZ#1344007)

CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY

Windows Server 2016 forest and domain functional levels now supported for trust

When using Identity Management, you can now establish a supported forest trust to Active Directory forests that run at the Windows Server 2016 forest and domain functional levels. (BZ#[1484683](#))

Directory Server no longer displays replication conflict entries in search results

Previously, if replication conflict entries existed in a replication topology, Directory Server returned them by default as part of the search result. As a consequence, certain LDAP clients behaved incorrectly if the server returned such entries. With this update, the server no longer returns conflict entries in a search and you have to explicitly request them. As a result, clients work as expected.

In addition, the update improves the resolution of more complex conflict scenarios.

For further details, see https://access.redhat.com/documentation/en-us/red_hat_directory_server/10/html/administration_guide/managing_replication-solving_common_replication_conflicts. (BZ#[1274430](#))

OpenLDAP is now compiled with OpenSSL instead of NSS

Previously, the OpenLDAP suite used the Mozilla implementation of Network Security Services (Mozilla NSS). With this update, OpenLDAP uses the OpenSSL library. Existing certificates in the NSS database (DB) are automatically extracted to the PEM format and passed to OpenSSL.

Note that NSS DBs continue to be supported. However, OpenSSL-like configuration, such as PEM files, is preferred over NSS-like configuration, such as NSS DB. (BZ#[1400578](#))

Samba rebased to version 4.7.1

The samba packages have been upgraded to upstream version 4.7.1, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Previously, the default value of the **rpc server dynamic port range** parameter was **1024-1300**. With this update, the default has been changed to **49152-65535** and now matches the range used in Windows Server 2008 and later. Update your firewall rules if necessary.
- Samba now uses the Advanced Encryption Standard (AES) instruction set of Intel CPUs to accelerate Server Message Block (SMB) 3 signing and encryption operations.
- The options of the **ntlm auth** parameter have been extended. The parameter now accepts the **ntlmv2-only** (alias **no**), **ntlmv1-permitted** (alias **yes**), **mschapv2-and-ntlmv2-only**, and **disabled** options. Additionally, the default value was renamed from **no** to **ntlmv2-only**.
- The **smbclient** utility no longer displays a banner with the domain, operating system, and server version when connecting to a server.
- The default value of the **client max protocol** parameter has been changed to **SMB3_11**. This enables utilities, such as **smbclient**, to connect to servers using the SMB 3.11 protocol without setting the protocol version.
- For a better interoperability, Samba no longer supports using mixed minor versions in a **CTDB** cluster.

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

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

- <https://www.samba.org/samba/history/samba-4.7.0.html> (BZ#1470048)

The SSSD LDAP provider can now automatically create user private groups for users

When using the System Security Services Daemon (SSSD) LDAP provider, a user group must be assigned to each user. Previously, the administrator had to create a group for each user manually. With this update, SSSD automatically generates a user private group from the user entry and ensures that the UID and GID match. To activate this feature, enable the **auto_private_groups** option in the LDAP provider section in the `/etc/sss/sss.conf` file. (BZ#1327705)

SSSD enrolled to an AD domain remembers the discovered AD site after the first successful connection

Previously, the System Security Services Daemon (SSSD) sent an LDAP ping to any Active Directory (AD) domain controller (DC) in order to determine a client's AD site. If the contacted DC was unreachable, a timeout occurred, which delayed the connection for several seconds. With this update, SSSD remembers the client's site after the first successful discovery. All subsequent LDAP pings are performed on the DC from the client's site, which helps speed up the request. (BZ#1400614)

SSSD logs changes in its status to syslog

Previously, the System Security Services Daemon (SSSD) logged information about changing its online or offline status to the SSSD logs only. With this update, the changes in SSSD status are logged also to the syslog service, which improves the availability of the information to system administrators. (BZ#1416150)

SSSD performance has improved

This update provides several performance-related enhancements for the System Security Services Daemon (SSSD). Most notably:

- Several missing indexes have been added in the SSSD cache, which makes lookups of cached objects faster.
- Changes to how users and groups are saved prevent the SSSD cache performance degradation that occurred after the cache was populated with a large number of cached objects.

As a result, SSSD reads user and group objects, especially large groups, faster. Also, the SSSD cache performance can now remain stable even when the cache size and the number of cache objects increase. (BZ#1472255, BZ#1482555)

The `pwdhash` utility can now retrieve the storage scheme from the configuration directory

Previously, if you passed the path to the configuration directory to the `pwdhash`, the utility used the default storage scheme of Directory Server to encrypt the password. With this update, the `pwdhash` utility uses the storage scheme set in the `nsslapd-rootpwstoragescheme` attribute in the `cn=config` entry, if you run `pwdhash` as a user with read permissions on the `/etc/dirsrv/slapd-instance_name/dse.ldif` file. As a result, you no longer have to specify the storage scheme in the mentioned scenario if it differs from the Directory Server's default. (BZ#1467777)

New utility to compare two Directory Server instances

This update adds the `ds-replcheck` utility to Directory Server. This utility compares the data of two servers in online mode, or two LDIF-formatted files in offline mode. As a result, you can now verify the replication consistency of two Directory Servers.

For further details, see https://access.redhat.com/documentation/en-us/red_hat_directory_server/10/html/administration_guide/comparing_two_directory_server_databases. (BZ#1406351)

Directory Server now supports enabling the `memberof` plug-in on read-only replicas

If you previously enabled the `memberof` plug-in on a read-only Directory Server replica server, the plug-in failed to update member entries. To use the plug-in in a replication topology, you could only enable it on write-enabled servers, and replicate the `memberof` attribute to read-only replicas. With this update, you can now alternatively enable the plug-in on all servers. As a result, you can use the plug-in on read-only servers the same as on write-enabled server.

For further details, see https://access.redhat.com/documentation/en-us/red_hat_directory_server/10/html/administration_guide/advanced_entry_management#considerations_vin. (BZ#1352121)

Directory Server rebased to version 1.3.7.5

The 389-ds-base packages have been upgraded to upstream version 1.3.7.5, 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-1-3-7-2.html>
- <http://directory.fedoraproject.org/docs/389ds/releases/release-1-3-7-3.html>
- <http://directory.fedoraproject.org/docs/389ds/releases/release-1-3-7-4.html>
- <http://directory.fedoraproject.org/docs/389ds/releases/release-1-3-7-5.html> (BZ#1470169)

Directory Server supports additional password storage schemes

For compatibility reasons, this update adds support for the following weak password storage schemes to Directory Server:

- CRYPT-MD5
- CRYPT-SHA256
- CRYPT-SHA512

For security reasons, use these weak storage schemes only temporary for existing installations and consider migrating to a strong password storage schema. (BZ#1479012)

Directory Server now uses separate normalized DN caches for each worker thread

Previously, multiple worker threads used a single normalized Distinguished Name (DN) cache. Consequently, if multiple clients performed operations on Directory Server, performance decreased. With this update, Directory Server now creates separate normalized DN caches for each worker thread. As a result, performance no longer decreases in the mentioned scenario. (BZ#1458536)

pkc-core rebased to version 10.5.1

The `pkc-core` packages have been upgraded to upstream version 10.5.1, which provides a number of bug fixes and enhancements over the previous version. Notably, this update addresses the requirements for the Common Criteria Protection Profile for Certification Authorities Version 2.1. (BZ#1473452)

Certificate System supports installing CA, KRA, and OCSP subsystems with CMC

This enhancement provides a mechanism to install CA, KRA, or OCSP subsystems with Certificate Management over CMS (CMC). The installation will be done in two steps. The first step of the installation will generate the Certificate Signing Requests (CSR) for the system certificates. The CSRs can be used

to issue the system certificates using CMC. The second step of the installation will use these system certificates and complete the subsystem installation. (BZ#1464549)

Certificate System supports creating instances running as a different user

Previously, Certificate System only used the systemd unit file from the `/usr/lib/systemd/system/` directory to start the service. Consequently, it was not possible to create a server running as a different user or group as `pkuser`. The `pkispawn` utility has been updated. If the configuration file passed to `pkispawn` contains a different user or group, the utility now creates an override file with the customized values in the `/etc/systemd/system/pki-tomcatd@<instance_name>.service.d/user.conf` file. As a result, running Certificate System user a different user or group as the default is possible. (BZ#1523410)

Certificate System can now create PKCS #12 files using PBES2 with PBKDF2 key derivation

This update enhances Certificate System and adds support for AES encryption of private keys recovered from the Key Recovery Authority (KRA), when token-based key recovery is disabled. Specifically, when AES encryption is enabled, exported PKCS #12 files containing the recovered key uses the PKCS #5 version 2.0 Password-Based Cryptography Specification version 2 (PBES2) with Password-Based Key Derivation Function 2 (PBKDF2) key derivation and AES 128 encryption. Using PBES2 with PBKDF2 makes the files created by Certificate System more secure. (BZ#1446786)

Certificate System CAs can now process CMC renewal requests signed by a previously issued signing certificate

This update enables the Certificate Authority (CA) to process Certificate Management over CMS (CMC) renewal requests signed by a previously issued signing certificate. The implementation uses the `caFullCMCUserSignedCert` with the `UniqueKeyConstraint` enhanced profile constraint, which has also been updated to disallow renewal of a key shared by a revoked certificate. Additionally, it preserves the `origNotAfter` attribute of the most recent certificate that shares the same key in the request, which allows the attribute to be used by the `RenewGracePeriodConstraint`. If there is an existing `origNotAfter` attribute, it is not overwritten in this process in order to not interfere with the existing `renewal by serial` flow. Additionally, the `caFullCMCUserSignedCert.cfg` profile has been updated to contain both the `UniqueKeyConstraint` and the `RenewGracePeriodConstraint`, which must be placed in the correct order. Note that by default, the `allowSameKeyRenewal` parameter is set to `true` in the `UniqueKeyConstraint`. (BZ#1419761)

Certificate System now uses the Mozilla NSS secure random number generator

With this update, Certificate System uses a secure random number generator provided by the Mozilla Network Security Services (NSS). This enables Red Hat Certificate System to synchronize its Deterministic Random Bit Generator (DRBG) with Red Hat Enterprise Linux as required by the Federal Information Processing Standard (FIPS) standard. (BZ#1452347)

Audit event changes in Certificate System

To provide more concise audit logs in Certificate System, the list of audit events enabled by default has been updated. Additionally, certain events have been merged or renamed.

For a full list of audit events in Red Hat Certificate System, including information in which subsystems they are enabled by default, see https://access.redhat.com/documentation/en-us/red_hat_certificate_system/9/html/administration_guide/audit_events. (BZ#1445532)

krb5 now includes the `kdcpolicy` interface

This update introduces the Kerberos key distribution center (KDC) policy interface, known as `kdcpolicy`, to the `krb5` package. Using `kdcpolicy`, administrators can provide a plug-in to `krb5`, which enables them to control ticket lifetimes and gives them more fine-grained control on service ticket issuance.

For details, see the MIT Kerberos Documentation: <https://web.mit.edu/kerberos/krb5-1.16/doc/plugindev/kdcpolicy.html>. (BZ#1462982)

Certificate System now supports configurable hashing algorithms for the SKI extension

Previously, Certificate System only supported the SHA1 hashing algorithm when generating the Subject Key Identifier (SKI) certificate extension. With this update, administrators can now configure the hashing algorithm for the SKI extension in certificate profiles.

The following algorithms are now available:

- SHA1
- SHA256
- SHA384
- SHA512

Note that the default algorithm is still SHA1. Therefore, existing profiles will not automatically be updated. (BZ#1024558)

The `pki` command-line interface automatically creates a default NSS database

The `pki` command-line interface requires a Network Security Services (NSS) database and its password to run operations over SSL connections, including basic authentication using a user name and password. Previously, `pki` displayed an error if the database did not exist or the database password was not specified. The command-line interface has been updated to automatically create a default NSS database without a password in the `~/ .dogtag/nssdb/` directory. As a result, operations over SSL can be executed without specifying an NSS database or password. (BZ#1400645)

Certificate System disables weak 3DES ciphers by default

By default, Certificate System now disables ciphers based on the weak Triple Data Encryption Standard (3DES). This increases the security of the system. However, administrators are able to enable these ciphers again, if needed. For details, see https://access.redhat.com/documentation/en-us/red_hat_certificate_system/9/html/administration_guide/configuring-ciphers.

As a result, new Certificate System installations have only strong ciphers enabled by default. (BZ#1469169)

The Certificate System CA subsystem's OCSP provider now includes the `nextUpdate` field in responses

If the Certificate Authority (CA) is configured to use the Certificate Revocation List (CRL) cache, the CA subsystem's Online Certificate Status Protocol (OCSP) responder now includes the `nextUpdate` field in OCSP responses. As a result, in such scenarios, clients which conform to the Lightweight OCSP Profile (RFC 5019) are now able to process OCSP responses. (BZ#1523443)

ding-libs rebased to version 0.6.1

The ding-libs packages have been upgraded to version 0.6.1. The most notable change is that ding-libs can now work with much bigger values, because the hard-coded limit to number of characters in values has been removed and the only limitation now is the amount of memory available. (BZ#1480270)

CHAPTER 6. CLUSTERING

New SNMP agent to query a Pacemaker cluster

The new **pcs_snmp_agent** agent allows you to query a Pacemaker cluster for data by means of SNMP. This agent provides basic information about a cluster, its nodes, and its resources. For information on configuring this agent, see the **pcs_snmp_agent**(8) man page and the High Availability Add-On Reference. (BZ#1367808)

Support for Red Hat Enterprise Linux High Availability clusters on Amazon Web Services

Red Hat Enterprise Linux 7.5 supports High Availability clusters of virtual machines (VMs) on Amazon Web Services (AWS). For information on configuring a Red Hat Enterprise Linux High Availability Cluster on AWS, see <https://access.redhat.com/articles/3354781>. (BZ#1451776)

Support for Red Hat Enterprise Linux High Availability clusters on Microsoft Azure

Red Hat Enterprise Linux 7.5 supports High Availability clusters of virtual machines (VMs) in Microsoft Azure. For information on configuring a Red Hat Enterprise Linux High Availability cluster on Microsoft Azure, see <https://access.redhat.com/articles/3252491>. (BZ#1476009)

Unfencing is done in resource cleanup only if relevant parameters changed

Previously, in a cluster that included a fence device that supports unfencing, such as **fence_scsi** or **fence_mpath**, a general resource cleanup or a cleanup of any stonith resource would always result in unfencing, including a restart of all resources. Now, unfencing is only done if the parameters to the device that supports unfencing changed. (BZ#1427648)

The pcsd port is now configurable

The port on which **pcsd** is listening can now be changed in the **pcsd** configuration file, and **pcs** can now communicate with **pcsd** using a custom port. This feature is primarily for the use of **pcsd** inside containers. (BZ#1415197)

Fencing and resource agents are now supported by AWS Python libraries and a CLI client

With this enhancement, Amazon Web Services Python libraries (python-boto3, python-botocore, and python-s3transfer) and a CLI client (awscli) have been added to support fencing and resource agents in high availability setups. (BZ#1512020)

Fencing in HA setups is now supported by Azure Python libraries

With this enhancement, Azure Python libraries (python-isodate, python-jwt, python-adal, python-msrest, python-msrestazure, and python-azure-sdk) have been added to support fencing in high availability setups. (BZ#1512021)

New features added to the sbd binary.

The **sbd** binary used as a command line tool now provides the following additional features:

- Easy verification of the functionality of a watchdog device
- Ability to query a list of available watchdog devices

For information on the **sbd** command line tool, see the **sbd**(8) man page. (BZ#1462002)

sbd rebased to version 1.3.1

The **sbd** package has been rebased to upstream version 1.3.1. This version brings the following changes:

- Adds commands to test and query watchdog devices
- Overhauls the command-line options and configuration file
- Properly handles **off** actions instead of **reboot** (BZ#1499864)

Cluster status now shows by default when a resource action is pending

Pacemaker supports a **record-pending** option that previously defaulted to **false**, meaning that cluster status would only show the current status of a resource (started or stopped). Now, **record-pending** defaults to **true**, meaning that cluster status may also show when a resource is in the process of starting or stopping. (BZ#1461976)

cluftr rebased to version 0.77.0

The cluftr packages have been upgraded to upstream version 0.77.0, which provides a number of bug fixes, new features, and user experience enhancements over the previous version. Among the notable updates are the following:

- When using **cluftr** to translate an existing configuration with the **pcs2pcscmd-needle** command in the case where the **corosync.conf** equivalent omits the **cluster_name** option (which is not the case with standard ``pcs``-initiated configurations), the contained **pcs cluster setup** invocation no longer causes cluster misconfiguration with the name of the first given node interpreted as the required cluster name specification. The same invocation will now include the **--encryption 0|1** switch when available, in order to reflect the original configuration accurately.
- In any script-like output sequence such as that produced with the **ccs2pcscmd** and **pcs2pcscmd** families of **cluftr** commands, the intended shell interpreter is now emitted in a valid form, so that the respective commented line can be honored by the operating system. (BZ#1381531)
- The **cluftr** tool now also covers some additional recently added means of configuration as facilitated with **pcs** (heuristics for a quorum device, meta attributes for top-level **bundle** resource units) when producing the sequence of configuring **pcs** commands to reflect existing configurations when applicable.

For information on the capabilities of **cluftr**, see the **cluftr(1)** man page or the output of the **cluftr -h** command. For examples of **cluftr** usage, see the following Red Hat Knowledgebase article: <https://access.redhat.com/articles/2810031>. (BZ#1509381)

Support for Sybase ASE failover

The Red Hat High Availability Add-On now provides support for Sybase ASE failover through the **ocf:heartbeat:sybaseASE** resource. To display the parameters you can configure for this resource, run the **pcs resource describe ocf:heartbeat:sybaseASE** command. For more information on this agent, see the **ocf_heartbeat_sybaseASE(7)** man page. (BZ#1436189)

CHAPTER 7. COMPILER AND TOOLS

The linuxptp package now supports active-backup bonding for clock synchronization

With this update, you can now specify a bond interface in the active-backup mode to be used by the **ptp4l** application. As a result, **ptp4l** uses the clock of the active interface from the bond as the **Precision Time Protocol (PTP)** clock and can switch to another interface of the bond in case of a failover. Additionally, the **phc2sys** utility in the automatic mode (the **-a** option) can synchronize the system clock to the **PTP** clock of the active interface when operating as a **PTP** slave and the **PTP** clock to the system clock when operating as a **PTP** master. (BZ#1002657)

parted can now resize partitions using the resizepart command

The ability to resize disk partitions using the **resizepart NUMBER END** command is now backported to the **parted** disk partitioning utility distributed with Red Hat Enterprise Linux 7. See the **parted(8)** man page for information.

Note that this command only resizes partitions, not file systems residing on them. Use file system utilities such as **resize2fs** to grow or shrink file systems. (BZ#1423357)

binutils rebased to version 2.27

The binutils package has been rebased to upstream version 2.27. This version brings the following changes:

- Support for compressed debug sections
- Improved handling of orphan sections during linking
- Support for the LLVM plugin
- Ability to insert new symbols into an object file with the **objcopy** utility
- Support for the IBM POWER9 architecture
- Support for the ARMv8.1 and ARMv8.2 instruction set extensions

Additionally, this update fixes the following bugs:

- Previously, the binutils package did not contain the **standards.info** documentation file that describes the GNU Coding Standard. This file has been added and is available through the **info** command again.
- Previously, the **ld** linker on the IBM Power Systems architecture stored intermediate data in the first object file specified by the linker command line. As a consequence, the linker terminated unexpectedly with a segmentation fault if that file was not used in the output and was discarded. The linker has been modified to directly store the data in the output file and skip the intermediate storage in the input file. As a result, linking no longer fails with a segmentation fault in the described situation. (BZ#1385959, BZ#1356856, BZ#1467390, BZ#1513014)

pcp rebased to version 3.12.2

The Performance Co-Pilot (PCP) application has been rebased to version 3.12.2, which includes many enhancements and bug fixes.

Collector systems updates:

- The following Performance Metric Domain Agents (PMDAs) have been updated: **perfevent**, containers and CGroups, MySQL slave metrics, Linux per-process metrics, and Linux kernel metrics for entropy, **slabinfo**, IPv6 sockets, and NFSD worker threads.
- New PMDAs are now available: Prometheus endpoint and HAProxy.
- Device Mapper statistics now expose an API.

Monitor systems updates:

- The derived metrics language has been extended for all monitors.
- The **pmchart** charting utility includes fixes for timezone and display bugs.
- The **pmlogconf** configuration utility automatically enables the **hotproc** metric logging and adds **atop** metrics. Performance is now more optimized.
- The **pcp-atop** monitoring utility recognizes the new **--hotproc** option. Several bugs have been fixed.
- The **pcp-pidstat** and **pcp-mpstat** monitoring utilities recognize several new output options.
- The **pmrep** reporting utility now supports Comma-separated Values (CSV) output compatible with the **sadf** tool. New utilities for exporting PCP metrics to various formats have also been added: **pcp2zabbix**, **pcp2xml**, **pcp2json**, and **pcp2elasticsearch**. (BZ#[1472153](#))

Improved DWARF 5 support in various tools

Support for the DWARF debugging format version 5 has been extended in the following tools:

- The **eu-readelf** tool from the elfutils package now recognizes all DWARF 5 tags and attributes.
- The **readelf** and **objdump** tools from the binutils package now recognize the DWARF 5 tag **DW_AT_exported_symbols** and correctly report its presence in debug information sections. (BZ#[1472955](#), BZ#[1472969](#))

systemtap rebased to version 3.2

The **SystemTap** utility has been updated to upstream version 3.2. Notable enhancements include:

- Support for extraction of matched regular expression has been added.
- Probe aliases for accepting input from the standard input have been added.
- Translator diagnostics have been improved.
- Support for the new **statx** system call has been added.
- A new string function **strpos()** for detecting substring position has been added to the stap language.

Additionally, this update fixes the following bugs:

- Previously, the statistics extractor functions **@min()** and **@max()** returned incorrect values. As a consequence, scripts relying on these functions did not work properly. The **@min()** and **@max()** functions have been fixed to return the correct maximum and minimum values. As a result, the affected scripts now work as expected.

- Previously, some kernel tracepoints were inconsistently listed with the **stap -L** command, even when they could not be probed. **SystemTap** has been fixed so that the listed and probeable tracepoint sets match again.
- The **netdev.receive** probe has been fixed and can collect data again.
- The example script **nettop.stp** affected by the broken **netdev.receive** probe again works as expected.

Note that the kernel version in Red Hat Enterprise Linux does not support extended Berkeley Packet Filter (eBPF), and consequently the related upstream **SystemTap** features are not available. (BZ#[1473722](#), BZ#1490862, BZ#[1506230](#), BZ#1485228, BZ#1518462)

valgrind rebased to version 3.13.0

The **valgrind** package has been upgraded to version 3.13.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes are:

- **Valgrind** has been extended in several ways to run large programs. The amount of memory usable by **Valgrind** has been increased to 128 GB. As a consequence, the **Memcheck** tool supports running applications that allocate up to approximately 60 GB. Additionally, **Valgrind** can now load executable files up to 1200 MB in size.
- The tools **Memcheck**, **Helgrind**, and **Massif** can now use a new execution tree (xtree) representation to report heap consumption of the analyzed applications.
- The symbol demangler has been updated to support the C++11 standard and the Rust programming language.
- Failures with long blocks of code using AVX2 instructions on the Intel and AMD 64-bit architecture have been fixed.
- The 64-bit **timebase** register of the PowerPC architecture is no longer modeled by **Valgrind** as only 32-bit.
- Support for the IBM Power Systems architecture has been extended to include the ISA 3.0B specification.
- An alternative implementation of Load-Linked and Store-Conditional instructions for the 64-bit ARM architecture has been added. The alternative implementation is enabled automatically when required. To enable it manually, use the **--sim-hints=fallback-llsc** option. (BZ#[1473725](#), BZ#1508148)

ncat rebased to version 7.50

The **ncat** utility, which is provided by the **nmap-ncat** package, has been rebased to upstream version 7.50. This provides a number of bug fixes and new features over the previous version. Notable changes include:

- Support for **SOCKS5** authentication has been added.
- The **-z** option for quickly checking the status of a port has been added.
- The **--no-shutdown** option now also works in connect mode, not only in listen mode. (BZ#1460249)

rsync rebased to version 3.1.2

The rsync packages have been upgraded to upstream version 3.1.2, which provides a number of bug fixes and enhancements over the previous version.

This update introduces the following output changes:

- The default output format of numbers has been changed to 3-digit groups, for example, **1, 234, 567**.
- The output of the **--progress** option has been changed; the following strings have been shortened: **xfer** to **xfr**, and **to-check** to **to-chk**.

Notable enhancements in this version include:

- I/O handling has been improved, which results in faster data transfers.
- New **--info** and **--debug** options have been added for more fine-grained output.
- The ability to synchronize nano-second modified times has been added.
- New options, **--usermap**, **--groupmap**, and **--chown**, have been added for manipulating file ownership during the copy operation.
- A new **--preallocate** option has been added. (BZ#[1432899](#))

tcpdump can now analyze virtio traffic

The **tcpdump** utility now supports the **virtio-vsock** communication device. This makes it possible for **tcpdump** to filter and analyze virtio communication between a hypervisor and a guest virtual machine. (BZ#[1464390](#))

vim now supports C++11 syntax highlighting

Syntax highlighting for C++ in the **Vim** text editor has been enhanced to support the C++11 standard. (BZ#[1267826](#))

vim now supports the blowfish2 encryption method

Support for the **blowfish2** encryption method has been added to the **Vim** text editor. This method provides stronger encryption than **blowfish**. To set the **blowfish2** encryption method, use the **:setlocal cm=blowfish2** command. Note that files encrypted with **blowfish2** are compatible between Red Hat Enterprise Linux 7 and Red Hat Enterprise Linux 6. (BZ#[1319760](#))

The IO::Socket::SSL Perl module now uses the system-wide CA certificate store by default

Previously, if a TLS application based on the **IO::Socket::SSL** Perl module did not provide an explicit path to a certificate authority (CA) certificate, no authority was known, and the peer's identity could not be verified. With this update, the module uses the system-wide CA certificate store by default. However, it is possible to disable any certificate store by passing the **undef** value to the **SSL_ca_file** option of the **IO::Socket::SSL->new()** constructor. (BZ#[1402588](#))

perl-DateTime-TimeZone rebased to version 1.70

The perl-DateTime-TimeZone package has been upgraded to upstream version 1.70, which provides a number of bug fixes and enhancements over the previous version. Notably:

- With this update, it is possible to install Bugzilla version 5, which requires a more recent version of perl-DateTime-TimeZone than the system provided previously.

- The Olson time zone database has been updated to version 2017b. Previously, applications written in the **Perl** language that use the **DateTime::TimeZone** module mishandled time zones that changed their specifications since version 2013h due to the outdated database.
- Using a local time zone from a tainted time zone identifier has been fixed. (BZ#[1241818](#), BZ#[1101251](#))

system-config-kdump now support selecting of either automated or manual kdump memory settings when fadump is performed

This update adds fadump memory reservation support into the system-config-kdump packages. As a result, users can now select either **Automated kdump memory settings** or **Manual settings** when **Firmware assisted dump** is selected. (BZ#1384943)

conman rebased to version 0.2.8

The conman packages have been upgraded to upstream version 0.2.8, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Scalability has been improved.
- **Coverity Scan** and **Clang** warnings have been fixed to improve stability.
- Arbitrary limit on the number of Intelligent Platform Management Interface (IPMI) Serial Over LAN (SOL) consoles has been fixed.
- The default value of the **loopback** setting has been changed to **ON** in the **conman.conf** file. (BZ#[1435840](#))

Support for the TFTP window size option has been implemented

With this update, support for the **window size** option according to RFC 7440 has been implemented in the Trivial File Transfer Protocol (TFTP) server and client. When the **window size** option is used, data blocks are sent in batches, which significantly improves throughput. (BZ#[1328827](#))

curl now supports disabling GSSAPI with SOCKS5

New **--socks5-basic** and **--socks5-gssapi** options for the **curl** utility and a corresponding option **CURLOPT_SOCKS5_AUTH** for the libcurl library have been introduced to control the authentication methods for SOCKS5 proxies. (BZ#1409208)

The rsync utility now copies files with their original nanosecond part of the time stamp

Previously, the **rsync** utility ignored the nanosecond part of the time stamp of files. As a consequence, the nanosecond time stamp of newly created files was always zero. With this update, the **rsync** utility recognizes the nanosecond part. As a result, the newly copied files keep their original nanosecond time stamp on systems that support it. (BZ#[1393543](#))

tcpdump rebased to version 4.9.2

The tcpdump package has been upgraded to upstream version 4.9.2, which provides a number of bug fixes (for almost 100 CVEs) and enhancements over the previous version. Notable changes include:

- A segmentation fault with OpenSSL 1.1 has been fixed and OpenSSL usage has been improved.
- The buffer overflow vulnerabilities have been fixed.
- The infinite loop vulnerabilities have been fixed.

- Many buffer over-read vulnerabilities have been fixed. (BZ#1490842)

oProfile support for Intel Xeon processor family extended

oProfile has been extended to support the Intel Xeon Phi™ Processor x200 and x205 Product Families. (BZ#1465354)

Support for Intel Xeon v4 uncore performance events in `libpfm`, `pcp`, and `papi`

This update adds support for Intel Xeon v4 uncore performance events to the **libpfm** performance monitoring library, the **pcp** tool, and the **papi** interface. (BZ#1474999)

Memory copying performance improved on IBM POWER architectures

Previously, the **memcpy()** function from the GNU C Library (**glibc**) used unaligned vector load and store instructions on 64-bit IBM POWER systems. Consequently, when **memcpy()** was used to access device memory on POWER9 systems, performance would suffer. The **memcpy()** function has been enhanced to use aligned memory access instructions, to provide better performance for applications regardless of the memory involved on POWER9, without affecting the performance on previous generations of the POWER architecture. (BZ#1498925)

TAI clock macro available

Previously, the kernel provided the **CLOCK_TAI** clock, but the **CLOCK_TAI** macro to access it was missing in the **glibc** header file **time.h**. The macro definition has been added to the header file. As a result, applications can now access the **CLOCK_TAI** kernel clock. (BZ#1448822)

Support for selective use of 4 KiB page tables on IBM z Systems

This update adds the option **--s390-pgste** to the **ld** linker from the binutils package to mark applications for the IBM z Systems architecture that require 4 KiB memory page tables on the lowest level. As a result, use of this feature can be restricted only to applications that need it, allowing optimal use of space by all applications on the system. Note that the **qemu** backend no longer forces 4 KiB lowest level page tables on all running applications. Make sure to specify the new option if your applications require them. (BZ#1485398)

More efficient `glibc` functions on IBM z Systems

Support for additional instructions of the IBM z Systems architecture has been added to the **glibc** library. As a result, programs compiled for this architecture can benefit from the increased performance of the **glibc** functions. (BZ#1375235)

The `ld` linker no longer incorrectly combines position-dependent and independent code

Previously, the **ld** linker combined object files on the IBM z Systems platform without considering whether they have been built for Position Independent Executable (PIE). Because PIE and non-PIE code cannot be combined, it was possible to create executable files that could not run. The linker has been extended to detect mixing of PIE and non-PIE code and produce an error message in this case. As a result, broken executable files can no longer be created this way. (BZ#1406430)

python-virtualenv rebased to 15.1.0

The python-virtualenv package has been upgraded to version 15.1.0, which provides a number of bug fixes and enhancements over the previous version. With this update, the following bundled packages have been upgraded: setuptools to version 28.0.0 and pip to version 9.0.1. (BZ#1461154)

python-urllib3 supports IP addresses in `subjectAltName`

The python-urllib3 package, a Python HTTP module with connection pooling and file POST abilities, now supports IP addresses in the **subjectAltName** (SAN) fields. (BZ#1434114)

Support for retpolines added to GCC

This update adds support for retpolines to GCC. Retpolines are a technique used by the kernel to reduce overhead of mitigating Spectre Variant 2 attacks described in CVE-2017-5715. (BZ#1535655)

Shenandoah garbage collector is now fully supported

The low pause time **Shenandoah** garbage collector for **OpenJDK**, previously available as a Technology Preview, is now fully supported on the Intel 64, AMD64, and 64-bit ARM architectures. **Shenandoah** performs concurrent evacuation which allows users to run with large heaps without long pause times. For more information, see <https://wiki.openjdk.java.net/display/shenandoah/Main>. (BZ#1578075)

CHAPTER 8. DESKTOP

GNOME Shell rebased to version 3.26

In Red Hat Enterprise Linux 7.5, **GNOME Shell** has been rebased to upstream version 3.26. Notable enhancements include:

- System search now provides results with an updated layout which makes them easier to read and shows more items at once. Additionally, it is now possible to search for system actions.
- The **Settings** application has a new layout.
- Various ways to insert emoji have been introduced for GNOME 3.26. This includes the Characters application and Polari, the GNOME IRC client.
- Display settings of GNOME have been redesigned.
- GNOME 3.26 no longer shows status icons in the bottom left part of the screen. GNOME Classic, which is the default session, now contains the **TopIcons** extension by default to provide the status tray functionality. Users of other session types than GNOME Classic can install the **TopIcons** extension manually.

For the full list of changes, see <https://help.gnome.org/misc/release-notes/3.26/> (BZ#1481381)

gnome-settings-daemon rebased to version 3.26

gnome-settings-daemon has been rebased to enable the Wayland display server protocol, more specifically, fractional monitor scaling. Instead of a single gnome-settings-daemon process, the user can now notice a collection of processes named gsd-* running in their sessions. (BZ#1481410)

libreoffice rebased to version 5.3

The LibreOffice office suite, has been upgraded to version 5.3, which includes a number of enhancements over the previous version:

- **LibreOffice** introduces a new LibreOffice UI, called **MUFFIN** (My User Friendly & Flexible INterface).
- The LibreOffice Writer contains a new **Go to Page** dialog to navigate in the text area.
- The LibreOffice Writer also introduces new table styles feature.
- A new Arrows toolbox has been added to **LibreOffice**.
- In Calc, number formatting and default cell styles have been improved.
- A new Template Selector was added to LibreOffice Impress

LibreOffice Base can no longer read **Firebird** 2.5 data. Embedded .odb files created in previous versions of **LibreOffice** are not compatible with this version.

For the full list of changes, see <https://wiki.documentfoundation.org/ReleaseNotes/5.3> (BZ#1474303)

GIMP rebased to version 2.8.22

GNU Image Manipulation Program (GIMP) version 2.8.22 includes the following significant bug fixes and enhancements:

Core:

- Saving to existing .xcf.bz and .xcf.gz files now truncates the files and no longer creates large files
- Text layer created by gimp-text-fontname respects border when resized

GUI:

- Drawing performance in single window mode, especially with pixmap themes, has been improved
- On Paint Dynamics editor dialog, the **y** axis is now indicates **Rate** instead **Flow**
- Pulsing progress bar in splash screen indicates unknown durations
- Gamut warning color for LC-MS display filter has been fixed
- Unbolding of bold font on edit has been fixed
- Accidental renaming of wrong adjacent item is now eliminated

Plug-ins:

- When importing PSD files, creating a wrong layer group structure is now eliminated
- Large images or large resolution no longer cause a crash in the PDF plug-in
- Parsing invalid PCX files is now stopped early and a subsequent segmentation fault is thus eliminated
- The Escape key can no longer close the Python console
- Filter **Edge Detect/Difference of Gaussians** returns empty image
- When printing, the images are composed onto a white background to prevent printing a black box instead of a transparent image
- Color vision deficiency display filters have been fixed to apply gamma correction directly
- Script-Fu regex match now returns proper character indexes for Unicode characters
- Script-Fu modulo for large numbers has been fixed

Updated Translations include: Basque, Brazilian Portuguese, Catalan, Chinese (PRC), Czech, Danish, Finnish, German, Greek, Hungarian, Icelandic, Italian, Kazakh, Norwegian, Polish, Portugese, Slovak, Slovenian, Scottish Gaelic and Spanish. (BZ#1210840)

Inkscape rebased to version 0.92.2

The rebased **Inkscape**, vector graphics software, provides a number of enhancements over the previous version, including the following:

- Mesh Gradients are now supported.
- Many SVG2 and CSS3 properties are now supported, for example, paint-order, mix-blend-mode. However, not all are available from the GUI.
- All objects are listed in the new Object dialog box from where you can select, label, hide, and lock any object.

- Selection sets make it possible to group objects together regardless of the document structure.
- Guides can now be locked to avoid accidental movement.
- Several new path effects have been added, among them Envelope/Perspective, Lattice Deformation, Mirror, and Rotate Copies.
- Several extensions have been added including a seamless pattern extension. In addition, many extensions have been updated or been given new features.
- A colorblindness simulation filter was added.
- The spray tool and measure tool have received several new features.
- The Pencil tool can create interactive smoothing for lines.
- BSplines are available for the Pen tool.
- Checkerboard background can be used to more easily see object transparencies. (BZ#[1480184](#))

webkitgtk4 rebased to version 2.16

The webkitgtk4 package has been upgraded to version 2.16, which provides a number of enhancements over the previous version. Notable enhancements include:

- To reduce memory consumption, hardware acceleration is now enabled on demand.
- webkitgtk4 contains a new WebKitSetting plug-in to set the hardware acceleration policy.
- CSS Grid Layout is enabled by default.
- Private browsing has been improved by adding a new API to create ephemeral web views.
- A new API has been provided to handle website data.
- Two new debugging tools are now available: memory sampler and resource usage overlay.
- GTK+ font settings are now honored.
- Theme rendering performance is improved when using GTK+ version 3.20 and higher. (BZ#[1476707](#))

qt5 rebased to version 5.9.2

The qt5 packages have been upgraded to upstream version 5.9.2, which provides a number of bug fixes and enhancements over the previous version. Notably, qt5 now contains:

- improved performance and stability
- long term support
- improved C++11 support - note that Qt 5.9 now requires C++11 compliant compiler
- Qt Quick Controls 2 - a new module with support for embedded devices (BZ#[1479097](#))

New package: qgnomeplatform

The **qGnomePlatform** Theme module is now included in Red Hat Enterprise Linux. In GNOME Desktop Environment, it makes applications created with **Qt 5** honor the current visual settings. (BZ#[1479351](#))

ModemManager rebased to version 1.6.8

The ModemManager package has been upgraded to upstream version 1.6.8 to support newer modem hardware. This provides a number enhancements over the previous version. Notably, the version of the **libqmi** library has been upgraded to 1.18.0 and the **libmbim** library to 1.14.2. In addition, the **usb_modeswitch** tool has been upgraded to 2.5.1 and the usb-modeswitch-data package to 20170806. (BZ#[1483051](#))

New packages: libsbios

Red Hat Enterprise Linux 7.5 now includes the libsbios packages to support flash Trusted Platform Module (TPM) and Synaptics Micro Systems Technology (MST) hubs. Libsbios is a library and utilities that can be used by client programs to get information from standard BIOS tables, such as the SMBIOS table. (BZ#1463329)

mutter rebased to version 3.26

The mutter package has been upgraded to version 3.26, which provides a number of bug fixes and enhancements over the previous version.

The most significant bug fixes include:

- Unexpected termination when respawning shortcut inhibitor dialog
- Unexpected termination during monitor configuration migration
- Multihead regressions in X11 session
- Screen rotation regressions
- Unexpected termination when reconnecting tablet device

The list of notable enhancements includes:

- Support for running headless
- Support for snap packages for sandboxed app IDs
- Support for `_NET_RESTACK_WINDOW` and `ConfigureRequest` siblings
- mutter now exports `_NET_NUMBER_OF_DESKTOPS`
- mutter now allows resizing of tiled windows
- Key bindings have been resolved with non-latin layouts
- Support for export tiling information to clients
- Monitor layout is now remembered across sessions (BZ#[1481386](#))

The `SANE_USB_WORKAROUND` environmental variable can make older scanners usable with USB3

Previously, Scanner Access Now Easy (SANE) was unable to communicate with certain older types of scanners when they were plugged into a USB3 port. This update introduces the **`SANE_USB_WORKAROUND`** environmental variable, which can be set to **1** to eliminate this problem. (BZ#[1458903](#))

The libyami package added for better video stream handling

With this update, the **libyami** package has been added to Red Hat Enterprise Linux 7 to improve video stream handling. In particular, the video stream is parsed and decoded with the help of hardware acceleration. (BZ#1456906)

netpbm rebased to version 10.79.00

The **netpbm** packages have been upgraded to version 10.79.00, which provides a large number of bug fixes and enhancements to multiple programs included in these packages. For detailed change log, see the `/usr/share/doc/netpbm/HISTORY` file. (BZ#1381122)

Red Hat Enterprise Linux 7.5 supports libva

Libva is an implementation for the Video Acceleration API (VA-API).

VA-API is an open-source library and API specification that provides access to graphics hardware acceleration capabilities for video processing. It consists of a main library and driver-specific acceleration back ends for each supported hardware vendor. (BZ#1456903)

GStreamer now supports mp3

An MPEG-2 Audio layer III decoder, more commonly known as **mp3**, has been added to **GStreamer**. The **mp3** support is available through the **mpeg123** library and the corresponding **GStreamer** plug-in.

The user can download the **mp3** plug-in using **GNOME Software** or using the codec installer in various **GStreamer** applications. (BZ#1481753)

GNOME control-center rebased to version 3.26

In Red Hat Enterprise Linux 7.5, control-center has been rebased to upstream version 3.26. Notable enhancements include:

- **Night Light** is a new feature that changes the color of your displays according to the time of day. The screen color follows the sunrise and sunset times for a given location, or can be set to a custom schedule. **Night Light** works with both **X11** and **Wayland** display server protocols.
- This update introduces a new layout to the **Settings** application. The grid of icons has been replaced by a sidebar, which allows switching between different areas. In addition, the **Settings** window is bigger and can be resized.
- GNOME's **Network** settings have been improved. **Wi-Fi** now has its own dedicated settings area and **Network** settings dialogs have been updated.
- GNOME's **Display** settings have been redesigned. The new design brings relevant settings to the forefront. With multiple displays connected, there is a row of buttons, which allows choosing the preferred use. The new **Display** settings include a preview version of a new scaling setting. This allows the size of what is shown on the screen to be adjusted to match the density (often expressed as PPI or DPI) of your display. Note that **Wayland** is recommended over **X11**, as per-display configuration is not supported on the latter.
- The user interface of three other areas of the **Settings** application has been redesigned: **Online Accounts**, **Printers**, and **Users**. (BZ#1481407)

New package: emacs-php-mode

This update adds the new **emacs-php-mode** package to Red Hat Enterprise Linux 7. **emacs-php-mode** provides PHP mode for the Emacs text editor thus enabling better PHP editing. (BZ#1266953)

Dutch keyboard layout provided

The installation of Red Hat Enterprise Linux in Dutch now provides an additional keyboard map that mimics the US International map used in the Windows OS. The new **latn1-pre.mim** keymap file

enables the user to utilize single keymap, diacritics, and thus type both in the English and Dutch language with ease. (BZ#1058510)

CHAPTER 9. FILE SYSTEMS

SMB 2 and SMB 3 now support DFS

Distributed File System (DFS), which was previously supported only with the Server Message Block (SMB) protocol version 1, is now also supported in SMB 2 and SMB 3.

With this update, you can now mount DFS shares using the SMB 2 and SMB 3 protocols. (BZ#1481303)

File system DAX now performs better when mapping a large amount of memory

Prior to this enhancement, the Direct Access (DAX) feature mapped only 4KiB entries into application address space. This had a negative performance impact on workloads that mapped large amounts of memory, because it increased Translation Lookaside Buffer (TLB) pressure. With this update, the kernel supports 2MiB Page Middle Directory (PMD) faults in persistent memory mappings. This significantly reduces TLB pressure, and file system DAX now performs better when mapping a large amount of memory. (BZ#1457572)

quotacheck is now faster on ext4

The **quotacheck** utility now directly scans **ext4** file system metadata instead of analyzing each individual file for occupied disk size. If the file system contains many files, quota initialization and quota check are now significantly faster. (BZ#1393849)

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#1626526)

CHAPTER 10. HARDWARE ENABLEMENT

Broadcom 5880 smart card readers with the updated firmware are now supported

This update includes the USB ID entries for the updated firmware version of the Broadcom 5880 smart card readers and Red Hat Enterprise Linux is now able to properly recognize and use these readers.

Note that users with the Broadcom 5880 smart card readers using older firmware versions should update the firmware. See the Support section at www.dell.com for more information about the updating process. (BZ#1435668)

fwupd now supports Synaptics MST hubs

Red Hat Enterprise Linux 7.5 adds a plug-in for Synaptics MST hubs to the **fwupd** utility. This plug-in enables you to flash firmware and query firmware information for this device. (BZ#1420913)

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#1462329)

Improved RT throttling mechanism

The current real-time throttling mechanism prevents the starvation of non-real-time tasks by CPU intensive real-time tasks. When a real-time run queue is throttled, it allows non-real-time tasks to run or if there are none, the CPU goes idle. To safely maximize CPU usage by decreasing the CPU idle time, the **RT_RUNTIME_GREED** scheduler feature has been implemented. When enabled, this feature checks if non-real-time tasks are starving before throttling the real-time task. As a result, the **RT_RUNTIME_GREED** scheduler option guarantees some run time on all CPUs for the non-real-time tasks, while keeping the real-time tasks running as much as possible. (BZ#1401061)

VMware Paravirtual RDMA Driver

This enhancement update adds VMware Paravirtual RDMA Driver to Red Hat Enterprise Linux. This feature allows VMware users to deploy and use Red Hat Enterprise Linux-based VMs with PVRDMA devices. (BZ#1454965)

opal-prd rebased to version 5.9

The **opal-prd** daemon, which handles hardware-specific recovery processes, has been rebased to version 5.9. This enhancement update includes the following important fixes and notable enhancements:

- flush after logging to **stdio** in debug mode
- fixes for memory leaks
- fix for **opal-prd** command line options
- fix for **occ_reset** call
- API comment regarding nanosleep ranges
- the **pnor** file is no longer passed while starting **opal-prd**
- on FSP system host, **pnor** access interface is disabled
- add support for runtime OCC load/start in ZZ

Users of **opal-prd** are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. (BZ#1456536)

libreswan now supports NIC offloading

This update of the libreswan packages introduces support for the network interface controller (NIC) offloading. **Libreswan** now automatically detects the NIC hardware offload support, and the **nic-offload=auto|yes|no** option has been added for manual setup of this feature. (BZ#1463062)

Trusted Computing Group TPM 2.0 System API library and management utilities available

The following packages, which handle the Trusted Computing Group's Trusted Platform Module (TPM) 2.0 hardware and which were previously available as a Technology Preview, are now fully supported:

- The tpm2-tss package adds the Intel implementation of the TPM 2.0 System API library. This library enables programs to interact with TPM 2.0 devices.
- The tpm2-tools package adds a set of utilities for management and utilization of TPM 2.0 devices from user space. (BZ#[1463097](#), BZ#1463100)

new packages: tpm2-abrmd

This update adds the tpm2-abrmd packages to Red Hat Enterprise Linux 7. The tpm2-abrmd packages provide a system service that implements the Trusted Platform Module (TPM) 2.0 Access Broker (TAB) and Resource Manager (RM) specification from the Trusted Computing Group. (BZ#1492466)

CHAPTER 11. INSTALLATION AND BOOTING

Assigning mount points to existing block devices is now possible in Kickstart installations

A new **mount** command is now available in Kickstart. This command assigns a mount point to a particular block device with a file system, and it can also reformat it if you specify the **--reformat** option.

The difference between **mount** and other storage-related commands like **autopart**, **part**, or **logvol** is that with **mount** you do not need to describe the entire storage configuration in the Kickstart file, you only need to make sure that the specified block devices exist on the system. However, if you want to create the storage configuration instead of using an existing one, and mount the various devices, then you must use the other storage configuration commands.

You can not use **mount** with the other storage configuration commands in the same Kickstart file. (BZ#[1450922](#))

The **livemedia-creator** utility now provides a sample Kickstart file for UEFI systems

The example Kickstart files provided with the **livemedia-creator** packages have been updated to support 32 and 64-bit UEFI systems. The files are located in the **/usr/share/lorax-version/** directory.

Note that **livemedia-creator** must be run on a UEFI system or virtual machine to build bootable UEFI disk images. (BZ#[1458937](#))

New option for the **network** Kickstart command binding the device configuration file to the device MAC address

You can now use the new **--bindto=mac** option with the **network** Kickstart command to use the **HWADDR** parameter (the MAC address) instead of the default **DEVICE** in the device's **ifcfg** file on the installed system. This will bind the device configuration to the MAC instead of the device name.

Note that the new **--bindto** option is independent of the **network --device** Kickstart option. It will be applied to the **ifcfg** file even if the device was specified in the Kickstart file using its name, **link**, or **bootif**. (BZ#[1328576](#))

New options for Kickstart **%packages** allow configuring **Yum** timeout and number of retries

This update adds two new options for the **%packages** section in Kickstart files:

- **--timeout=X** - sets the **Yum** timeout to **X** seconds. Defaults to 30.
- **--retries=Y** - sets the number of **Yum** retries to **Y**. Defaults to 10.

Note that if you use multiple **%packages** sections during the installation, options set on the section which appears last will be used for every section. If the last section has neither of these options set, every **%packages** section in the Kickstart file will use the default values.

These new options may help when performing a large number parallel installations from a single package source at once, when package download speed is limited by disk read or network speeds. The new options only affect the system during installation and have no effect on **Yum** configuration on the installed system. (BZ#[1448459](#))

The Red Hat Enterprise Linux 7 ISO image can be used to create guests virtual machines on IBM z Systems

With this release, you can create a bootable Red Hat Enterprise Linux ISO file for KVM virtual machines on the IBM z Systems architecture. As a result, Red Hat Enterprise Linux guest virtual machines on IBM z Systems can boot from a **boot.iso** file. (BZ#[1478448](#))

ARPUPDATE option for ifcfg-* files has been introduced

This update introduces the ARPUPDATE option for the **ifcfg-*** files with default value **yes**. Setting the value to **no** allows administrators to disable updating neighboring computers with address resolution protocol (ARP) information about current network interface controller (NIC). This is especially needed when using Linux Virtual Server (LVS) Load Balancing with Direct routing enabled. (BZ#[1478419](#))

The --noconfig option added for the rpm -v command

With this update, the **--noconfig** option has been added to the **rpm -v** command. This option enables the command to list only the altered non-configuration files, which helps diagnose system problems. (BZ#[1406611](#))

ifcfg-* files now allow you to specify a third DNS server

ifcfg-* configuration files now support the **DNS3** option. You can use this option to specify a third Domain Name Server (DNS) address to be used in **/etc/resolv.conf**, instead of the previous maximum of two DNS servers. (BZ#[1357658](#))

Multi-threaded xz compression in rpm-build

This update adds multi-threaded **xz** compression for source and binary packages when setting the **%_source_payload** or **%_binary_payload** macros to the **wLTX.xzdio** pattern. In it, **L** represents the compression level, which is 6 by default, and **X** is the number of threads to be used (may be multiple digits), for example **w6T12.xzdio**. To enable this feature, edit the **/usr/lib/rpm/macros** file or declare the macro within the spec file or at the command line.

As a result, compressions take less time for highly parallel builds, which is beneficial especially for continuous integration of large projects that are built on hardware with many cores. (BZ#[1278924](#))

CHAPTER 12. KERNEL

Memory Protection Keys are now supported in later Intel processors

Memory Protection Keys provide a mechanism for enforcing page-based protections, but without requiring modifications of the page tables when an application changes protection domains. To determine if your processor supports Memory Protection Keys, check for the **pku** flag in the `/proc/cpuinfo` file. Further documentation including programming examples can be found in the `/usr/share/doc/kernel-doc-*/Documentation/x86/protection-keys.txt` file, which is provided by the kernel-doc package. (BZ#1272615)

EDAC support added for Pondicherry 2 memory controllers

Error Detection and Correction support has been added for Pondicherry 2 memory controllers used on machines based on the Intel Atom C3000-series processors. (BZ#1273769)

MBA is now supported

Memory Bandwidth Allocation (MBA) is an extension of the existing Cache QoS Enforcement (CQE) feature found in Broadwell servers. **MBA** is a feature of the Intel Resource Director Technology (RDT) that provides control over memory bandwidth for applications. With this update, the **MBA** support is added. (BZ#1379551)

Swap optimizations enable fast block devices to be used as secondary memory

Previously, the swap subsystem was not performance-critical because the performance of rotating disks, especially in terms of latency, was orders of magnitude worse than the rest of the memory management subsystem. With the advent of fast SSD devices, the overhead of the swap subsystem has become significant. This update brings a series of performance optimizations that reduce this overhead. (BZ#1400689)

HID Wacom rebased to version 4.12

The **HID Wacom** kernel module packages have been upgraded to upstream version 4.12, which provides a number of bug fixes and enhancements over the previous version:

- The **hid_wacom** power supply code has been updated, fixing previously existing problems.
- Support has been added for the Bluetooth-based Intuos 2 Pro pen tablet.
- Bugs affecting the Intuos 2 Pro pen tablet and the Bamboo slate have been fixed. (BZ#1475409)

New livepatch functionality improves the latency and success rate of the kpatch-patch packages

With this update, the **kpatch** kernel live patching infrastructure has been upgraded to use the new upstream **livepatch** functionality for patching the kernel. This functionality improves the scheduling latency and success rate of the kpatch-patch hotfix packages. (BZ#1430637)

Persistent Kernel Module Upgrade (PKMU) supported

The **kmod** packages provide various programs for automatic loading, unloading, and management of kernel modules. Previously, **kmod** searched for the modules only in the `/lib/modules/<kernel version>` directory. Consequently, users needed to perform additional actions, for example, run the `/usr/sbin/weak-modules` script to install symlinks, to make the modules loadable. With this update, **kmod** have been modified to search for the modules anywhere in the file system. As a result, users can now install new modules to a separate directory, configure the **kmod** tools to look for modules there, and the modules will be available automatically for the new kernel. Users can also specify several directories for a kernel, or different directories for different kernels. The kernel version is specified with a regular expression. (BZ#1361857)

The Linux kernel now supports encrypted SMB 3 connections

Prior to introducing this feature, the kernel only supported unencrypted connections when using the Server Message Block (SMB) protocol. This update adds encryption support for SMB 3.0 and later protocol versions. As a result, users can mount SMB shares using encryption, if the server provides or requires this feature.

To mount a share using the encrypted SMB protocol, pass the **seal** mount option together with the **vers** mount option set to **3.0** or later to the **mount** command. For further details and an example, see the **seal** parameter description in https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/storage_administration_guide/mounting_an_smb_share#tab.frequently_ (BZ#1429710)

SME enabled on AMD Naples platforms

With this update, AMD Secure Memory Encryption (SME) is provided by systems based on AMD Naples platforms. The Advanced Encryption Standard (AES) engine has the ability to encrypt and decrypt dynamic random access memory (DRAM). **SME**, provided by the AES engine, is intended to protect machines against hardware-probing attacks. To activate **SME**, boot the system with the kernel parameter **mem_encrypt=on**. (BZ#1361287)

Support for the ie31200_edac driver

This enhancement adds support for the ie31200_edac driver to the consumer version of Skylake and Kabi Lake CPU families. (BZ#1482253)

EDAC now supports GHES

This enhancement adds Error Detection and Correction (EDAC) support for using the Generic Hardware Error Source (GHES) provided by BIOS. GHES is now used as a source for memory corrected and uncorrected errors instead of a hardware specific driver. (BZ#1451916)

CUIR enhanced scope detection is now fully supported

Support for Control Unit Initiated Reconfiguration (CUIR) enables the Direct Access Storage Device (DASD) device driver to automatically take paths to DASDs offline for concurrent services. If other paths to the DASD are available, the DASD stays operational.

CUIR informs the DASD device driver when the paths are available again, and the device driver attempts to vary them back online.

In addition to the support for Linux instances running in Logical Partitioning (LPAR) mode, support for Linux instances on IBM z/VM systems has been added. (BZ#1494476)

kdump allows a vmcore collection without the root file system being mounted

In Red Hat Enterprise Linux 7.4, **kdump** required the root file system to be mounted although this is not always necessary for the collection of a **vmcore** image file. Consequently, **kdump** failed to collect a **vmcore** file if the root device could not be mounted when the dump target was not on the root file system, but, for example, on a usb or on the network. With this enhancement, if the root device is not required for dump, it is not mounted, and a **vmcore** file can be collected. (BZ#1431974, BZ#1460652)

KASLR fully supported and enabled by default

Kernel address space layout randomization (KASLR), which was previously available as a Technology Preview, is fully supported in Red Hat Enterprise Linux 7.5 on the AMD64 and Intel 64 architectures. KASLR is a kernel feature that contains two parts, kernel text KASLR and **mm** KASLR. These two parts work together to enhance the security of the Linux kernel.

The physical address and virtual address of kernel text itself are randomized to a different position separately. The physical address of the kernel can be anywhere under 64TB, while the virtual address of the kernel is restricted between [0xffffffff80000000, 0xffffffffc0000000], the 1GB space.

The starting address of three `mm` sections (the direct mapping, `vmalloc`, and `vmemmap` section) is randomized in a specific area. Previously, starting addresses of these sections were fixed values.

KASLR can thus prevent inserting and redirecting the execution of the kernel to a malicious code if this code relies on knowing where symbols of interest are located in the kernel address space.

KASLR code is now compiled in the Linux kernel, and it is enabled by default. If you want to disable it explicitly, add the `nokaslr` kernel option to the kernel command line. (BZ#1491226)

Intel® Omni-Path Architecture (OPA) Host Software

Intel® Omni-Path Architecture (OPA) host software is fully supported in Red Hat Enterprise Linux 7.5. 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_5_RN_J98644.pdf. (BZ#1543995)

`noreplace-paravirt` has been removed from the kernel command line parameters

The `noreplace-paravirt` kernel command line parameter has been removed, because the parameter is no longer compatible with the patches to mitigate the Spectre and Meltdown vulnerabilities. Booting AMD64 and Intel 64 systems with `noreplace-paravirt` in kernel command line will cause repeated reboots of the operating system. (BZ#1538911)

The new EFI `memmap` implementation is now available on SGI UV2+ systems

Prior to this update, the Extensible Firmware Interface (EFI) stable runtime services mapping across kexec reboot (`memmap`) implementation was not available on Silicon Graphics International (SGI) UV2 and later systems. This update adds support for EFI `memmap`. Additionally, this update also enables use of Secure Boot with the `kdump` kernel. (BZ#1102454)

Mounting pNFS shares with flexible file layout is now fully supported

Flexible file layout on pNFS clients was first introduced in Red Hat Enterprise Linux 7.2 as a Technology Preview. With Red Hat Enterprise Linux 7.5, it is now fully supported.

pNFS flexible file layout enables advanced features such as non-disruptive file mobility and client-side mirroring, which provides enhanced usability in areas such as databases, big data, and virtualization. See <https://datatracker.ietf.org/doc/draft-ietf-nfsv4-flex-files/> for detailed information about pNFS flexible file layout. (BZ#1349668)

CHAPTER 13. NETWORKING

Error handling in the output of the `dhcp-script` has been improved

Previously, any error in the output of the `dhcp-script` was ignored. With this update the output of the script is logged on the `add`, `old`, `del`, `arp-add`, `arp-del`, `tftp` actions. As a result, errors are displayed while `dnsmasq` is running.

Note that the lease-init action happens only at a start of `Dnsmasq`. With this update, only a summary of the output is logged and not the standard error output, which passes to the `systemd` service for logging. (BZ#[1188259](#))

Network namespace isolation has been added to `ipset`

Previously, `ipset` entries were visible and could be modified by any network namespace. This update provides `ipset` with isolation per network namespace. As a result, `ipset` configuration is separated for each namespace. (BZ#[1226051](#))

NetworkManager now supports multiple routing tables to enable source routing

This update adds a new `table` attribute for IPv4 and IPv6 routes which can be configured manually by the user. For each manual static route, a routing table can be selected. As a result, configuring the table of a route has the effect of configuring the route in that table. Additionally, the default routing table of a connection profile can be configured via the new `ipv4.route-table` and `ipv6.route-table` settings for IPv4 and IPv6 respectively. These settings determine in which table the routes are placed, except manual routes that explicitly overwrite this setting. (BZ#[1436531](#))

nftables rebased to version 0.8

The nftables packages have been upgraded to version 0.8, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- New expressions: `fib`, `numgen`, `quota`, `rt`, `notrack` have been added.
- Support hashing of any arbitrary key combination has been added.
- Support to set non-byte bound packet header fields, including checksum adjustment has been added.
- Variable reference for set element definitions and variable definitions from element commands can now be used.
- Support to flush set has been added.
- Support for logging flags has been added.
- Support for `tc classid` parser has been added.
- Endianness problems with link layer address have been solved.
- Parser to keep map flag around on definition has been fixed.
- The time datatype now uses milliseconds, as the kernel expects. (BZ#[1472261](#))

Persistent DHCP client behavior added to NetworkManager

With this update, the `ipv4.dhcp-timeout` property can be set to either the `maximum for a 32-bit integer (MAXINT32)` value or to the `infinity` value. As a result, `NetworkManager` never stops trying to get or renew a lease from a `DHCP` server until it is successful. (BZ#[1350830](#))

NetworkManager exposes new properties to expose team options

Previously, **NetworkManager** applied team configuration to connections providing a JSON string to the **config** property, which was the only property available in the team setting. This update adds new properties in **NetworkManager** matching one to one the team configuration options. As a result, the configuration may be provided either through a unique JSON string in the **NetworkManager config** property or assigning values to the new team properties. Any configuration change applied in **config** is reflected to the new team properties and vice versa. The correct configuration of team link-watchers and team.runner is now enforced in **NetworkManager**. Wrong or unknown link-watcher and team.runner configurations result in the full team connection being rejected.

Note that when changing the brand new **runner** property, all the properties related to specific runners are reset to default. (BZ#1398925)

Packets mark is now reflected on replies

Previously, when receiving a connection request on a closed port, an error packet was sent back to the client. When the incoming connection was marked with some firewall rules, the generated error message did not have this mark because this functionality was not implemented in the kernel. With this update, the generated error message has the same marking as the incoming packet that tried to initiate the connection. (BZ#1469857)

New Socket timestamping options for NTP

This update adds the **SOF_TIMESTAMPING_OPT_PKTINFO** and **SOF_TIMESTAMPING_OPT_TX_SWHW** socket timestamping options for hardware timestamping with bonding and other virtual interfaces in **Network Time Protocol (NTP)** implementations, such as chrony. (BZ#1421164)

iproute2 rebased to version 4.11.0

The iproute2 package has been upgraded to upstream version 4.11.0, which provides a number of bug fixes and enhancements. Notably, the **ip** tool includes:

- Support for JSON output to various commands has been added.
- Support for more interface type attributes has been added.
- Support for colored output has been added.
- Support for the **label**, **dev** options and the **rule** objects in **ip-monitor** state.
- Support for selectors in the **ip-rule** command has been added.

Additionally, notable improvements for the **tc** utility include:

- Support for the bash-completion function for **tc**.
- The **vlan** action in **tc** has been introduced.
- The extended mode in the **pedit** action has been introduced.
- Stream Control Transmission Protocol (SCTP) support in the **csum** action has been added.

For other tools:

- Support for extended statistics in the **lstat** tool has been added.
- Support for **SCTP** in the **nstat** utility has been added. (BZ#1435647)

The **tc-pedit** action now supports offset relative to Layer 2 and Layer 4

The **tc-pedit** action allows modification of packet data. This update adds support for specifying the **offset** options relative to the Layer 2, 3 and 4 headers to **tc-pedit**. This makes **pedit header** handling more robust and flexible. As a result, editing Ethernet header is more convenient and accessing the Layer 4 header works independently to the Layer 3 header size. (BZ#1468280)

Features backported to **iproute**

A number of enhancements have been backported to the **iproute** package. Notable changes include:

- Pipeline debug support has been added to the **devlink** tool via the **dpipe** subcommand.
- Hardware offload status is now available in the **tc** filter, indicated by the **in_hw** or **not_in_hw** flags.
- Support for IPv6 in the **tc pedit** action has been added.
- Setting and retrieving **eswitch** encapsulation support has been added to the **devlink** tool.
- Matching capabilities of the **tc flower** filter have been enhanced:
- Support for matching on TCP flags.
- Support for matching on the type-of-service (ToS) and the time-to-live (TTL) fields in the IP header.

(BZ#1456539)

The **Geneve** driver rebased to version 4.12

The **Geneve** driver has been updated to version 4.12, which provides several bug fixes and enhancements for Open vSwitch (OVS) or Open Virtual Network (OVN) deployments using **Geneve** tunneling. (BZ#1467288)

A control switch added for **VXLAN** and **GENEVE** offloading

This update adds a new control switch to the **ethtool** utility to enable or disable offloading of the **VXLAN** and **GENEVE** tunnels to network cards. This enhancement enables easier debugging of issues with the **VXLAN** or **GENEVE** tunnels. In addition, you can resolve issues caused by offloading these types of tunnels to network cards by using **ethtool** to disable the feature. (BZ#1308630)

unbound rebased to version 1.6.6

The **unbound** packages have been rebased to upstream version 1.6.6, which provides a number of bug fixes and enhancements over the previous version. Notable changes are as follows:

- DNS Query Name (QNAME) minimisation according to RFC 7816 has been implemented.
- A new **max-udp-size** configuration option has been added; its default value is **4096**.
- A new **DNS64** module and a new **dns64-prefix** option have been added.
- New **insecure_add** and **insecure_remove** commands have been added to the **unbound-control** utility for administration of negative trust anchors.
- The **unbound-control** utility is now capable of bulk addition and removal of local zones and local data. To perform these actions, use the **local_zones**, **local_zones_remove**, **local_datas**, and **local_datas_remove** commands.

- The **libldns** is no longer a dependency of **libunbound** and will not be installed with it.
- A new **so-reuseport**: option is now available for distributing queries evenly over threads on Linux.
- New Resource Record types have been added: **CDS**, **CDNSKEY**, **URI** (according to RFC 7553), **CSYNC**, and **OPENPGPKEY**.
- New **local-zone** types have been added: **inform** to log a message with a client IP and **inform_deny** to log a query and drop the answer to it.
- Remote control over local sockets is now available; use the **control-interface**: **/path/sock** and **control-use-cert**: **no** commands.
- A new **ip-transparent**: configuration option has been added for binding to non-local IP addresses.
- A new **ip-freebind**: configuration option has been added for binding to an IP address while the interface or address is down.
- A new **harden-algo-downgrade**: configuration option has been added.
- The following domains are now blocked by default: **onion** (according to RFC 7686), **test**, and **invalid** (according to RFC 6761).
- A user-defined pluggable event API for the **libunbound** library has been added.
- To set the working directory for **Unbound**, either use the **directory**: **dir** with the **include**: **file** statement in the **unbound.conf** file, which ensures that the includes are relative to the directory, or use the **chroot** command with an absolute path.
- Fine-grained localzone control has been implemented with the following options: **define-tag**, **access-control-tag**, **access-control-tag-action**, **access-control-tag-data**, **local-zone-tag**, and **local-zone-override**.
- A new **outgoing-interface**: **netblock/64** IPv6 option has been added to use Linux freebind feature for every query with a random 64-bit local part.
- Logging of DNS replies has been added, which is similar to query logs.
- Trust anchor signaling has been implemented that uses key tag query and **trustanchor.unbound CH TXT** queries.
- Extension mechanisms for DNS (EDNS) Client subnet has been implemented.
- **ipsecmod**, an opportunistic IPsec support module, has been implemented. (BZ#1251440)

DHCP now supports standard dynamic DNS updates

With this update, the DHCP server allows updating DNS records by using a standard protocol. As a result, DHCP supports standard dynamic DNS updates as described in RFC 2136:

<https://tools.ietf.org/html/rfc2136>. (BZ#1394727)

DDNS now supports additional algorithms

Previously, the **dhcpcd** daemon supported only the **HMAC-MD5** hashing algorithm which is considered insecure for critical applications. As a consequence, the **Dynamic DNS (DDNS)** updates were

potentially insecure. This update adds support for additional algorithms: **HMAC-SHA1**, **HMAC-SHA224**, **HMAC-SHA256**, **HMAC-SHA384**, or **HMAC-SHA512**. (BZ#1396985)

IPTABLES_SYSCTL_LOAD_LIST now supports the sysctl.d files

The **sysctl** settings in **IPTABLES_SYSCTL_LOAD_LIST** are reloaded by the **iptables** init script when the **iptables** service is restarted. The modified settings were previously searched only in the **/etc/sysctl.conf** file. This update adds support for searching these modifications in the **/etc/sysctl.d/** directory as well. As a result, the user-provided files in **/etc/sysctl.d/** are now correctly taken into account when the **iptables** service is restarted. (BZ#1402021)

SCTP now supports MSG_MORE

The **MSG_MORE** flag is set to buffer small pieces of data until a full packet is ready for transmission or until a call is performed that does not specify this flag. This update adds support for **MSG_MORE** on the Stream Control Transmission Protocol (SCTP). As a result, small data chunks can be buffered and sent as a full packet. (BZ#1409365)

MACsec rebased to version 4.13

The **Media Access Control Security (MACsec)** driver has been upgraded to upstream version 4.13, which provides a number of bug fixes and enhancements over the previous version. Notable enhancements include:

- **Generic Receive Offload (GRO)** and **Receive Packet Steering (RPS)** are enabled on **MACsec** devices.
- The **MODULE_ALIAS_GENL_FAMILY** module has been added. This helps tools such as **wpa_supplicant** to start even if the module is not loaded yet. (BZ#1467335)

Enhanced performance when using the m1x5 driver in Open vSwitch

The Open vSwitch (OVS) application enables Virtual Machines to communicate with each other and the physical network. OVS resides in the hypervisor and switching is based on twelve tuple matching on flows. However, the OVS software-based solution is very CPU-intensive. This affects the system performance and prevents using the fully available bandwidth.

With this update, the **m1x5** driver for Mellanox ConnectX-4, ConnectX-4 Lx, and ConnectX-5 adapters can offload OVS. The Mellanox Accelerated Switching And Packet Processing (ASAP2) Direct technology enables offloading OVS by handling the OVS data-plane in Mellanox ConnectX-4 and later network interface cards with Mellanox Embedded Switch or eSwitch, while maintaining an unmodified OVS control-plane. As a result, the OVS performance is significantly higher and less CPU-intensive.

The current actions supported by ASAP2 Direct include packet parsing and matching, forward, drop along with VLAN push/pop, or VXLAN encapsulation and decapsulation. (BZ#1456687)

The Netronome NFP Ethernet driver now supports the representor netdev feature

This update backports the **representor netdev** feature for the Netronome NFP Ethernet driver to Red Hat Enterprise Linux 7.5. This enhancement enables the driver:

- To receive and transmit fallback traffic
- To be used in Open vSwitch
- To support programming flows to the NFP hardware by using the TC-Flower utility (BZ#1454745)

Support for offloading TC-Flower actions

This update adds support for offloading the **TC-Flower** classifier and actions related to offloading of Open vSwitch. This allows acceleration of Open vSwitch using Netronome SmartNICs. (BZ#1468286)

DNS stub resolver improvements

The DNS stub resolver in the **glibc** package has been updated to the upstream glibc version 2.26. Notable improvements and bug fixes include:

- Changes to the **/etc/resolv.conf** file are now automatically recognized and applied to running programs. To restore the previous behavior, add the **no-reload** option to the **options** line in **/etc/resolv.conf**. Note that depending on system configuration, the **/etc/resolv.conf** file might be automatically overwritten as part of the configuration of the networking subsystem, removing the **no-reload** option.
- The previous limit of six search domain entries is removed. You can now specify any number of domains with the **search** directive in **/etc/resolv.conf**. Note that additional entries may add significant overhead to DNS processing; consider running a local caching resolver if the number of entries exceeds three.
- The handling of various boundary conditions in the **getaddrinfo()** function is fixed. Very long lines in the **/etc/hosts** file (including comments) no longer affect lookup results from other lines. Unexpected terminations related to stack exhaustion on systems with certain **/etc/hosts** configuration no longer occur.
- Previously, when the **rotate** option was enabled in **/etc/resolv.conf**, the first DNS query of a new process was always sent to the second name server configured in the name server list in **/etc/resolv.conf**. This behavior has been changed, and the first DNS query now randomly selects a name server from the list. Subsequent queries rotate through the available name servers, as before. (BZ#677316, BZ#1432085, BZ#1257639, BZ#1452034, BZ#1329674)

CHAPTER 14. SECURITY

LUKS-encrypted removable storage devices can be now automatically unlocked using NBDE

With this update, the **clevis** package and the **clevis-udisks2** subpackage enable users to bind removable volumes to a Network-Bound Disk Encryption (NBDE) policy. To automatically unlock a LUKS-encrypted removable storage device, such as a USB drive, use the **clevis luks bind** and **clevis luks unlock** commands. (BZ#1475408)

new package: clevis-systemd

This update of the **Clevis** pluggable framework introduces the **clevis-systemd** subpackage, which enables administrators to set automated unlocking of LUKS-encrypted non-root volumes at boot time. (BZ#1475406)

openSCAP can be now integrated into Ansible workflows

With this update, the **OpenSCAP** scanner can generate remediation scripts in the form of Ansible Playbooks, either based on profiles or based on scan results. Playbooks based on SCAP Security Guide Profiles contain fixes for all rules, and playbooks based on scan results contain only fixes for rules that fail during an evaluation. The user can also generate a playbook from a tailored Profile, or customize it directly by editing the values in the playbook. Tags, such as Rule ID, strategy, complexity, disruption, or references, used as metadata for tasks in playbooks serve to filter, which tasks to apply. (BZ#1404429)

SECCOMP_FILTER_FLAG_TSYNC enables synchronization of calling process threads

This update introduces the **SECCOMP_FILTER_FLAG_TSYNC** flag. When adding a new filter, this flag synchronizes all other threads of the calling process to the same **seccomp** filter tree. See the **seccomp(2)** man page for more information.

Note that if an application installs multiple **libseccomp** or **seccomp-bpf** filters, the **seccomp()** syscall should be added to the list of allowed system calls. (BZ#1458278)

nss rebased to version 3.34

The nss packages have been upgraded to upstream version 3.34, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- TLS compression is no longer supported.
- The TLS server code now supports session ticket without an RSA key.
- Certificates can be specified using a PKCS#11 URI.
- The **RSA-PSS** cryptographic signature scheme is now allowed for signing and verification of certificate signatures. (BZ#1457789)

SSLv3 disabled in mod_ssl

To improve the security of SSL/TLS connections, the default configuration of the **httpd mod_ssl** module has been changed to disable support for the **SSLv3** protocol, and to restrict the use of certain cryptographic cipher suites. This change will affect only fresh installations of the **mod_ssl** package, so existing users should manually change the SSL configuration as required.

Any SSL clients attempting to establish connections using **SSLv3**, or using a cipher suite based on **DES** or **RC4**, will be denied in the new default configuration. To allow such insecure connections, modify the **SSLProtocol** and **SSLCipherSuite** directives in the **/etc/httpd/conf.d/ssl.conf** file. (BZ#1274890)

Libreswan now supports split-DNS configuration for IKEv2

This update of the libreswan packages introduces support for split-DNS configuration for the Internet Key Exchange version 2 (IKEv2) protocol through the **leftmodecfgdns=** and **leftcfgdomains=** options. This enables the user to reconfigure a locally running DNS server with DNS forwarding for specific private domains. (BZ#1300763)

libreswan now supports AES-GMAC for ESP

With this update, support for Advanced Encryption Standard (AES) Galois Message Authentication Code (GMAC) within IPsec Encapsulating Security Payload (ESP) through the **phase2alg=null_auth_aes_gmac** option has been added to the libreswan packages. (BZ#1475434)

openssl-ibmca rebased to 1.4.0

The openssl-ibmca packages have been upgraded to upstream version 1.4.0, which provides a number of bug fixes and enhancements over the previous version:

- Added Advanced Encryption Standard Galois/Counter Mode (AES-GCM) support.
- Fixes for **OpenSSL** operating in FIPS mode incorporated. (BZ#1456516)

opencryptoki rebased to 3.7.0

The opencryptoki packages have been upgraded to upstream version 3.7.0, which provides a number of bug fixes and enhancements over the previous version:

- Upgraded the license to Common Public License Version 1.0 (CPL).
- Added ECDSA with SHA-2 support for Enterprise PKCS #11 (EP11) and Common Cryptographic Architecture (CCA).
- Improved performance by moving from mutex locks to Transactional Memory (TM). (BZ#1456520)

atomic scan with configuration_compliance enables creating security-compliant container images at build time

The **rhel7/openscap** container image now provides the **configuration_compliance** scan type. When used as an argument for the **atomic scan** command, this new scan type enables users to:

- scan Red Hat Enterprise Linux-based container images and containers against any profile provided by the SCAP Security Guide (SSG)
- remediate Red Hat Enterprise Linux-based container images to be compliant with any profile provided by the SSG
- generate an HTML report from a scan or a remediation.

The remediation results in a container image with an altered configuration that is added as a new layer on top of the original container image.

Note that the original container image remains unchanged and only a new layer is created on top of it. The remediation process builds a new container image that contains all the configuration improvements. The content of this layer is defined by the security policy of scanning. This also means that the remediated container image is no longer signed by Red Hat, which is expected, since it differs from the original container image by containing the remediated layer. (BZ#1472499)

tang-nagios enables Nagios to monitor Tang

The **tang-nagios** subpackage provides the **Nagios** plugin for **Tang**. The plugin enables the **Nagios** program to monitor a **Tang** server. The subpackage is available in the Optional channel. See the **tang-nagios(1)** man page for more information. (BZ#1478895)

clevis now logs privileged operations

With this update, the **clevis-udisks2** subpackage logs all attempted key recoveries to the Audit log, and the privileged operations can be now tracked using the Linux Audit system. (BZ#1478888)

PK11_CreateManagedGenericObject() has been added to NSS to prevent memory leaks in applications

The **PK11_DestroyGenericObject()** function does not destroy objects allocated by **PK11_CreateGenericObject()** properly, but some applications depend on a function for creating objects that persist after the use of the object. For this reason, the **Network Security Services** (NSS) libraries now include the **PK11_CreateManagedGenericObject()** function. If you create objects with **PK11_CreateManagedGenericObject()**, the **PK11_DestroyGenericObject()** function also properly destroys underlying associated objects. Applications, such as the **curl** utility, can now use **PK11_CreateManagedGenericObject()** to prevent memory leaks. (BZ#1395803)

openssh now supports openssl-ibmca and openssl-ibmpkcs11 HSMs

With this update, the **OpenSSH** suite enables hardware security modules (HSM) handled by the **openssl-ibmca** and **openssl-ibmpkcs11** packages. Prior to this, the **OpenSSH** seccomp filter prevented these cards working with the **OpenSSH** privilege separation. The seccomp filter has been updated to allow system calls needed by the cryptographic cards on IBM z Systems. (BZ#1478035)

cgroup_seclabel enables fine-grained access control on cgroups

This update introduces the **cgroup_seclabel** policy capability that enables users to set labels on control group (cgroup) files. Prior to this addition, labeling of the cgroup file system was not possible, and to run the **systemd** service manager in a container, read and write permissions for any content on the cgroup file system had to be allowed. The **cgroup_seclabel** policy capability enables fine-grained access control on the cgroup file system. (BZ#1494179)

The boot process can now unlock encrypted devices connected by network

Previously, the boot process attempted to unlock block devices connected by network before starting network services. Because the network was not activated, it was not possible to connect and decrypt these devices.

With this update, the **remote-cryptsetup.target** unit and other patches have been added to **systemd** packages. As a result, it is now possible to unlock encrypted block devices that are connected by network during system boot and to mount file systems on such block devices.

To ensure correct ordering between services during system boot, you must mark the network device with the **_netdev** option in the **/etc/crypttab** configuration file.

A common use case for this feature is together with network-bound disk encryption. For more information on network-bound disk encryption, see the following chapter in the Red Hat Enterprise Linux Security Guide:

https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/sec-using_network-bound_disk_encryption (BZ#1384014)

SELinux now supports InfiniBand object labeling

This release introduces **SELinux** support for **InfiniBand** end port and P_Key labeling, including enhancements to the kernel, policy, and the **semanage** tool. To manage **InfiniBand**-related labels, use the following commands:

- **semanage ibendport**
- **semanage ibpkey** (BZ#[1471809](#), BZ#[1464484](#), BZ#[1464478](#))

libica rebased to 3.2.0

The libica packages have been upgraded to upstream version 3.2.0, which most notably adds support for the Enhanced SIMD instructions set. (BZ#[1376836](#))

SELinux now supports systemd No New Privileges

This update introduces the **nnp_nosuid_transition** policy capability that enables SELinux domain transitions under **No New Privileges** (NNP) or **nosuid** if **nnp_nosuid_transition** is allowed between the old and new contexts. The selinux-policy packages now contain a policy for **systemd** services that use the **NNP** security feature.

The following rule describes allowing this capability for a service:

```
allow source_domain target_type:process2 { nnp_transition
nosuid_transition };
```

For example:

```
allow init_t fprintd_t:process2 { nnp_transition nosuid_transition };
```

The distribution policy now also contains the m4 macro interface, which can be used in SELinux security policies for services that use the **init_nnp_daemon_domain()** function. (BZ#[1480518](#))

Libreswan rebased to version 3.23

The libreswan packages have been upgraded to upstream version 3.23, which provides a number of bug fixes, speed improvements, and enhancements over the previous version. Notable changes include:

- Support for the extended DNS Security Extensions (DNSSEC) suite through the **dnssec-enable=yes|no**, **dnssec-rootkey-file=**, and **dnssec-anchors=** options.
- Experimental support for Postquantum Preshared Keys (PPK) through the **ppk=yes|no|insist** option.
- Support for Signature Authentication (RFC 7427) for RSA-SHA.
- The new **logip=** option with the default value **yes** can be used to disable logging of incoming IP addresses. This is useful for large-scale service providers concerned for privacy.
- Unbound DNS server **ipsecmod module** support for Opportunistic IPsec using **IPSECKEY** records in DNS.
- Support for the Differentiated Services Code Point (DSCP) architecture through the **decap-dscp=yes** option. DSCP was formerly known as Terms Of Service (TOS).
- Support for disabling Path MTU Discovery (PMTUD) through the **nopmtudisc=yes** option.
- Support for the IDr (Identification - Responder) payload for improved multi-domain deployments.
- Resending IKE packets on extremely busy servers that return the **EAGAIN** error message.
- Various improvements to the updown scripts for customizations.

- Updated preferences of crypto algorithms as per RFC 8221 and RFC 8247.
- Added the **%none** and **/dev/null** values to the **leftupdown=** option for disabling the updown script.
- Improved support for rekeying using the CREATE_CHILD_SA exchange.
- IKEv1 XAUTH thread race conditions resolved.
- Significant performance increase due to optimized pthread locking.

See the **ipsec.conf** man page for more information. (BZ#[1457904](#))

libreswan now supports IKEv2 MOBIKE

This update introduces support for the IKEv2 Mobility and Multihoming (MOBIKE) protocol (RFC 4555) using the XFRM_MIGRATE mechanism through the **mobike=yes|no** option. MOBIKE enables seamless switching of networks, for example, Wi-Fi, LTE, and so on, without disturbing the IPsec tunnel. (BZ#[1471763](#))

scap-workbench rebased to version 1.1.6

The scap-workbench packages have been upgraded to version 1.1.6, which provides a number of bug fixes and enhancements over the previous version. Notable changes are:

- Added support for generating Bash and Ansible remediation roles from profiles and for scanning results. The generated remediations can be saved to a file for later use.
- Added support for opening tailoring files directly from the command line.
- Fixed a short integer overflow when using SSH port numbers higher than 32,768. (BZ#[1479036](#))

OpenSCAP is now able to generate results for DISA STIG Viewer

The **OpenSCAP** suite is now able to generate results in the format compatible with the **DISA STIG Viewer** tool. This enables the user to scan a local system for Defense Information Systems Agency Security Technical Implementation Guide (DISA STIG) compliance and open results in **DISA STIG Viewer**. (BZ#[1505517](#))

selinux-policy no longer contains permissive domains

As a security hardening measure, the **SELinux** policy now does not set the following domains to permissive mode by default:

- blkmapd_t
- hsqldb_t
- ipmievtd_t
- sanlk_reseted_t
- systemd_hwdb_t
- targetd_t

The default mode for these domains is now set to enforcing. (BZ#[1494172](#))

audit rebased to version 2.8.1

The audit packages have been upgraded to upstream version 2.8.1, which provides a number of bug fixes and enhancements over the previous version. Notable changes are:

- Added support for ambient capability fields.
- The **Audit** daemon now works also on IPv6.
- Added the default port to the **auditd.conf** file.
- Fixed the **auvirt** tool to report Access Vector Cache (AVC) messages. (BZ#[1476406](#))

opensc now supports the SCE7.0 144KDI CAC Alt. tokens

This update adds support for the SCE7.0 144KDI Common Access Card (CAC) Alternate tokens. These new cards were not compliant with the previous U.S. Department of Defense (DoD) Implementation Guide for CAC PIV End-Point specification, and the **OpenSC** driver has been updated to reflect the updated specification. (BZ#[1473418](#))

CHAPTER 15. SERVERS AND SERVICES

Leftover dbus processes

Red Hat Enterprise Linux 7.5 adds a feature that enables users to launch **dbus**-using applications remotely, for example over SSH or over IBM Platform LSF.

However, when processes using **dbus** are launched remotely, **dbus** processes keep running even after the main process is closed, blocking the remote session and preventing it from terminating properly.

To work around this problem, follow the instructions at <https://access.redhat.com/solutions/3257651>. (BZ#1460262)

dbus rebased to version 1.10

The dbus packages have been upgraded to upstream version 1.10, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- **dbus-run-session** is a new utility to run a **dbus** session bus for the runtime of a login session, making **ssh** sessions which start dbus-using applications more predictable and reliable. See **man 1 dbus-run-session** for more details.
- Several memory and file descriptor leaks have been fixed. This improves the **dbus-daemon** memory usage and reliability.
- The well-known system and session bus configuration files have been moved from **/etc/dbus-1/** to the **/usr/share/dbus-1/** directory. While the old location can still be used, it is deprecated (specifically, **session.conf** and **system.conf** are deprecated, but system administrator configuration snippets under **session.d** and **system.d** are permitted). (BZ#1480264)

tuned rebased to version 2.9.0

The tuned packages have been upgraded to upstream version 2.9.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes include the following:

- The **net** plug-in has been extended with the **ring** and **pause** parameters.
- The concept of manually or automatically set profile has been introduced.
- A directory for profile recommendation files is now supported. (BZ#1467576)

chrony rebased to version 3.2

The chrony packages have been upgraded to upstream version 3.2, which provides a number of bug fixes and enhancements over the previous version. Notable enhancements include:

- Support for hardware timestamping with bonding, bridging, and other logical interfaces that aggregate ethernet interfaces
- Support for transmit-only hardware timestamping with network cards that can timestamp only received Precision Time Protocol (PTP) packets but not Network Time Protocol (NTP) packets
- Improved stability of synchronization with hardware timestamping and interleaved modes
- An improved **leapsectz** option to automatically set the offset of the system clock between International Atomic Time (TAI) and Coordinated Universal Time (UTC) (BZ#1482565)

SNMP page counting can be now disabled in CUPS

The simple network management protocol (SNMP) page counting currently shows incorrect information for certain printers. With this update, the CUPS printing system supports turning off the SNMP page counting, which prevents the problem. To do so, add ***cupsSNMPPages: False** into the printer's postscript printer description (PPD) file.

The procedure for adding options into printer's PPD file is described in solution article at <https://access.redhat.com/solutions/1427573> . (BZ#1434153)

CUPS can be set to use only ciphers from TLS version 1.2 or later

The CUPS printing system can now be set to use only ciphers from TLS version 1.2 or later. You can use the functionality by adding the option **SSLOptions MinTLS1.2** into the **/etc/cups/client.conf** file for the CUPS client or into the **/etc/cups/cupsd.conf** file for the CUPS daemon. (BZ#1466497)

The squid packages now provide the kerberos_ldap_group helper

This update adds the **kerberos_ldap_group** external Access Control Lists (ACL) helper to the squid packages. The **kerberos_ldap_group** helper is a reference implementation that supports Simple Authentication and Security Layer (SASL) and Generic Security Services API (GSSAPI) authentication to an LDAP server, intended primarily to connect to Active Directory or OpenLDAP-based LDAP servers. (BZ#1452200)

OpenIPMI rebased to version 2.0.23

The OpenIPMI packages have been upgraded to version 2.0.23, which provides a number of bug fixes and enhancements. Among others:

- It adds a command to set a duty cycle of the fans directly.
- It adds a way to specify the state directory from the command line after the compilation time.
- It changes the message map size to 32 bits so that it can handle a full 16-message window.
- It adds support for the IPMI LAN Simulator commands. See the `ipmi_sim_cmd(5)` man page.
- It adds support for the IPMI LAN Interface configuration file. See the `ipmi_lan(5)` man page. (BZ#1457805)

Overview of changes from freeIPMI 1.2.9 to freeIPMI 1.5.7

These are the most important changes:

- The **ipmi-fru** tool now supports the output of the DDR3 and DDR4 SDRAM modules and new FRU multirecords. - The new **ipmi-config** tool is a consolidated configuration tool implementing all the functionalities that were previously in the **bmc-config**, **ipmi-pef-config**, **ipmi-sensors-config**, and **ipmi-chassis-config** tools. - The **ipmi-sel** tool reads and manages the IPMI System Event Log records, which makes the tool useful for debugging the system.

A complete list of changes is available after the installation in the **/usr/share/doc/freeipmi/NEWS** file. (BZ#1435848)

A new clear_env option available in PHP FPM pool configuration

This update introduces a new **clear_env** option in PHP's FastCGI Process Manager (FPM) pool configuration. If the **clear_env** option is disabled, environment variables set when running the FPM daemon are preserved and available to scripts. By default, **clear_env** is enabled, preserving current behavior. (BZ#1410010)

CHAPTER 16. STORAGE

Data Deduplication and Compression with VDO

Red Hat Enterprise Linux 7.5 introduces Virtual Data Optimizer (VDO). This feature enables you to create block devices that transparently provide data deduplication, compression, and thin provisioning. Standard file systems and applications can run on these virtual block devices without modification.

VDO is currently available only on the AMD64 and Intel 64 architectures.

For more information on VDO, see the chapter Data Deduplication and Compression with VDO in the Storage Administration Guide: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/storage_administration_guide/vdo. (BZ#1480047)

New boom utility for managing LVM snapshot and image boot entries

This release adds the **boom** command, which you can use to manage additional boot loader entries on the system. You can use it to create, delete, list, and modify auxiliary boot entries for system snapshots and images. The utility provides a single tool for managing boot menu entries for LVM snapshots; therefore you no longer need to manually edit boot loader configuration files and work with detailed kernel parameters. The tool is provided by the `lvm2-python-boom` package. (BZ#1278192)

DM Multipath no longer requires reservation keys in advance

DM Multipath now supports two new configuration options in the `multipath.conf` file:

- **unpriv_sgio**
- **prkeys_file**

The **reservation_key** option of the **defaults** and **multipaths** sections accepts a new keyword: **file**. When set, the **multipathd** service will now use the file configured in the **prkeys_file** option of the **defaults** section to get the reservation key to use for the paths of a multipath device. The **prkeys** file is automatically updated by the **mpathpersist** utility. The default for the **reservation_key** option remains undefined, and default for the **prkeys_file** is `/etc/multipath/prkeys`.

If the new **unpriv_sgio** option is set to **yes**, DM Multipath will now create all new devices and their paths with the **unpriv_sgio** attribute. This option is used internally by other software, and is unnecessary for most DM Multipath users. It defaults to **no**.

These changes make it possible to use the **mpathpersist** utility without knowing ahead of time what reservation keys will be used and without adding them to the `multipath.conf` configuration file. As a result, it is now easier to use the **mpathpersist** utility to manage multipath persistent reservations in multiple setups. (BZ#1452210)

New property parameter supported in **blacklist** and **blacklist_exception** sections of `multipath.conf`

The `multipath.conf` configuration file now supports the **property** parameter in the **blacklist** and **blacklist_exception** sections of the file. This parameter allows users to blacklist certain types of devices. The **property** parameter takes a regular expression string that is matched against the **udev** environment variable names for the device.

The **property** parameter in **blacklist_exception** works differently than the other **blacklist_exception** parameters. If the parameter is set, the device must have a **udev** variable that matches. Otherwise, the device is blacklisted.

Most usefully, this parameter allows users to blacklist SCSI devices that multipath should ignore, such as

USB sticks and local hard drives. To allow only SCSI devices that could reasonably be multipathed, set this parameter to (**SCSI_IDENT_|ID_WWN**) in the **blacklist_exceptions** section of the **multipath.conf** file. (BZ#1456955)

smartmontools now support NVMe devices

This update adds support for Nonvolatile Memory Express (NVMe) devices, especially Solid-state Drive (SSD) disks, into the smartmontools package. As a result, the smartmontools utilities can now be used for monitoring NVMe disks with the Self-Monitoring, Analysis and Reporting Technology System (S.M.A.R.T.). (BZ#1369731)

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

SCSI T10 DIF/DIX is fully supported in Red Hat Enterprise Linux 7.5, provided that the hardware vendor has qualified it and provides full support for the particular HBA and storage array configuration. DIF/DIX is not supported on other configurations, it is not supported for use on the boot device, and it is not supported on virtualized guests.

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

FUJITSU supports DIF and DIX on:

EMULEX 16G FC HBA:

- EMULEX LPe16000/LPe16002, 10.2.254.0 BIOS, 10.4.255.23 FW, with:
- FUJITSU ETERNUS DX100 S3, DX200 S3, DX500 S3, DX600 S3, DX8100 S3, DX8700 S3, DX8900 S3, DX200F, DX60 S3, AF250, AF650, DX60 S4, DX100 S4, DX200 S4, DX500 S4, DX600 S4, AF250 S2, AF650 S2

QLOGIC 16G FC HBA:

- QLOGIC QLE2670/QLE2672, 3.28 BIOS, 8.00.00 FW, with:
- FUJITSU ETERNUS DX100 S3, DX200 S3, DX500 S3, DX600 S3, DX8100 S3, DX8700 S3, DX8900 S3, DX200F, DX60 S3, AF250, AF650, DX60 S4, DX100 S4, DX200 S4, DX500 S4, DX600 S4, AF250 S2, AF650 S2

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

EMC supports DIF on:

EMULEX 8G FC HBA:

- LPe12000-E and LPe12002-E with firmware 2.01a10 or later, with:
- EMC VMAX3 Series with Enginuity 5977; EMC Symmetrix VMAX Series with Enginuity 5876.82.57 and later

EMULEX 16G FC HBA:

- LPe16000B-E and LPe16002B-E with firmware 10.0.803.25 or later, with:
- EMC VMAX3 Series with Enginuity 5977; EMC Symmetrix VMAX Series with Enginuity 5876.82.57 and later

QLOGIC 16G FC HBA:

- QLE2670-E-SP and QLE2672-E-SP, with:
- EMC VMAX3 Series with Enginuity 5977; EMC Symmetrix VMAX Series with Enginuity 5876.82.57 and later

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

Support for DIF/DIX remains in Technology Preview for other HBAs and storage arrays. (BZ#1499059)

File system Direct Access (DAX) and device DAX now support huge pages

Previously, each file system DAX and device DAX page fault mapped to a single page in the user space. With this update, file system DAX and device DAX can now map persistent memory in larger chunks, called huge pages.

File system DAX supports huge pages that are, for example, 2 MiB in size on the AMD64 and Intel 64 architectures, and device DAX supports using either 2 MiB or 1 GiB huge pages on AMD64 and Intel 64. In comparison, a standard page is 4 KiB in size on these architectures.

When creating a DAX namespace, you can configure the page size that the namespace should use for all page faults.

Huge pages lead to fewer page faults, smaller page tables, and less Translation Lookaside Buffer (TLB) contention. As a result, file system DAX and device DAX now use less memory and perform better. (BZ#1457561, BZ#1383493)

fsadm can now grow and shrink LUKS-encrypted LVM volumes

The **fsadm** utility is now able to grow and shrink Logical Volume Manager (LVM) volumes that are encrypted with Linux Unified Key Setup (LUKS). This applies both to using **fsadm** directly with the **fsadm --lvresize** command and to using it indirectly through the **lvresize --resizefs** command.

Note that due to technical limitations, resizing of encrypted devices with a detached header is not supported. (BZ#1113681)

CHAPTER 17. SYSTEM AND SUBSCRIPTION MANAGEMENT

cockpit rebased to version 154

The cockpit packages, which provide the **Cockpit** browser-based administration console, have been upgraded to version 154. This version provides a number of bug fixes and enhancements. Notable changes include:

- The **Accounts** page now enables the configuration of account locking and password expiry.
- Load graphs consistently ignore loopback traffic on all networks.
- **Cockpit** provides information about unmet conditions for **systemd** services.
- Newly created timers on the **Services** page are now started and enabled automatically.
- It is possible to dynamically resize the terminal window to use all available space.
- Various navigation and JavaScript errors with Internet Explorer have been fixed.
- **Cockpit** uses Self-Signed Certificate Generator (SSCG) to generate SSL certificates, if available.
- Loading SSH keys from arbitrary paths is now supported.
- Absent or invalid **/etc/os-release** files are now handled gracefully.
- Unprivileged users now cannot use the shutdown/reboot button on the **System** page.

Note that certain cockpit packages are available in the Red Hat Enterprise Linux 7 Extras channel; see <https://access.redhat.com/support/policy/updates/extras>. (BZ#1470780, BZ#1425887, BZ#1493756)

Users of yum-utils now can perform actions prior to transactions

A new **yum-plugin-pre-transaction-actions** plug-in has been added to the **yum-utils** collection. It allows users to perform actions before a transaction starts. The usage and configuration of the plug-in are almost identical to the existing **yum-plugin-post-transaction-actions** plug-in. (BZ#1470647)

yum can disable creation of per-user cache as a non-root user

New **usercache** option has been added to the **yum.conf(5)** configuration file of the **yum** utility. It allows the users to disable the creation of per-user cache when **yum** runs as a non-root user. The reason for this change is that in some cases users do not want to create and populate per-user cache, for example in cases where the space in the **\$TMPDIR** directory is consumed by the user cache data. (BZ#1432319)

yum-builddep now allows to define RPM macros

The **yum-builddep** utility has been enhanced to allow you to define RPM macros for a **.spec** file parsing. This change has been made because, in some cases, RPM macros need to be defined in order for **yum-builddep** to successfully parse a **.spec** file. Similarly to the **rpm** utility, the **yum-builddep** tool now allows you to specify RPM macros with the **--define** option. (BZ#1437636)

subscription-manager now displays the host name upon registration

Until now, the user needed to search for the effective host name for a given system, which is determined by different Satellite settings. With this update, the **subscription-manager** utility displays the host name upon the registration of the system. (BZ#1463325)

A subscription-manager plugin now runs with yum-config-manager

With this update, the **subscription-manager** plugin runs with the **yum-config-manager** utility. The **yum-config-manager** operations now trigger **redhat.repo** generation, allowing Red Hat Enterprise Linux containers to enable or disable repositories without first running **yum** commands. (BZ#[1329349](#))

subscription-manager now protects all product certificates in /etc/pki/product-default/

Previously, the **subscription-manager** utility only protected those product certificates provided by the **redhat-release** package whose tag matched **rhel-#**. Consequently, product certificates such as **RHEL-ALT** or **High Touch Beta** were sometimes removed from the **/etc/pki/product-default/** directory by the **product-id yum** plugin. With this update, **subscription-manager** has been modified to protect all certificates in **/etc/pki/product-default/** against automatic removal. (BZ#[1526622](#))

rhnmigrateclassic-to-rhsm now automatically enables the subscription-manager and product-id yum plugins

With this update, the **rhnmigrateclassic-to-rhsm** utility automatically enables the **yum** plugins: **subscription-manager** and **product-id**. With this update, the **subscription-manager** utility automatically enables the **yum** plugins: **subscription-manager** and **product-id**. This update benefits users of Red Hat Enterprise Linux who previously used the **rhncclienttools** utility to register their systems to Red Hat Network Classic or who still use it with Satellite 5 entitlement servers, and who have temporarily disabled the **yum** plugins. As a result, **rhnmigrateclassic-to-rhsm** allows an easy transition to using the newer **subscription-manager** tools for entitlements. Note that running **rhnmigrateclassic-to-rhsm** displays a warning message indicating how to change this default behavior if it is not desirable. (BZ#[1466453](#))

subscription-manager now automatically enables the subscription-manager and product-id yum plugins

With this update, the **subscription-manager** utility automatically enables the **yum** plugins: **subscription-manager** and **product-id**. This update benefits users of Red Hat Enterprise Linux who previously used the **rhncclienttools** utility to register their systems to Red Hat Network Classic or who still use it with Satellite 5 entitlement servers, and who have temporarily disabled the **yum** plugins. As a result, it is easier for users to start using the newer **subscription-manager** tools for entitlements. Note that running **subscription-manager** displays a warning message indicating how to change this default behavior if it is not desirable. (BZ#[1319927](#))

subscription-manager-cockpit replaces subscription functionality in cockpit-system

This update introduces a new **subscription-manager-cockpit** RPM. The new **subscription-manager-cockpit** RPM provides a new dbus-based implementation and a few fixes to the same subscriptions functionality provided by **cockpit-system**. If both RPMs are installed, the implementation from **subscription-manager-cockpit** is used. (BZ#[1499977](#))

virt-who logs where the host-guest mapping is sent

The **virt-who** utility now uses the **rhsm.log** file to log the owner or account to which the host-guest mapping is sent. This helps proper configuration of **virt-who**. (BZ#[1408556](#))

virt-who now provides configuration error information

The **virt-who** utility now checks for common **virt-who** configuration errors and outputs log messages that specify the configuration items that caused these errors. As a result, it is easier for a user to correct **virt-who** configuration errors. (BZ#[1436617](#))

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)

CHAPTER 18. VIRTUALIZATION

KVM virtualization on IBM z Systems

KVM virtualization is now supported on IBM z Systems. However, this feature is only available in the newly introduced user space based on kernel version 4.14, provided by the kernel-alt packages.

Also note that due to hardware differences, certain features and functionalities of KVM virtualization differ from what is supported on AMD64 and Intel 64 systems.

For details on installing and using KVM virtualization on IBM z Systems, see the Virtualization Deployment and Administration Guide. (BZ#[1400070](#), BZ#1379517, BZ#1479525, BZ#1479526, BZ#1471761)

KVM virtualization supported on IBM POWER9

With this update, KVM virtualization is supported on IBM POWER9 systems, which makes it possible to use KVM virtualization on IBM POWER9 machines. However, this feature is only available in the newly introduced user space based on kernel version 4.14, provided by the kernel-alt packages.

Also note that due to hardware differences, certain features and functionalities of KVM virtualization on IBM POWER9 differ from what is supported on AMD64 and Intel 64 systems.

For details on installing and using KVM virtualization on POWER9 systems, see the Virtualization Deployment and Administration Guide. (BZ#[1465503](#), BZ#1478482, BZ#1478478)

KVM virtualization supported on IBM POWER8

With this update, KVM virtualization is supported on IBM POWER8 systems, which makes it possible to use KVM virtualization on IBM POWER8 machines.

Note that due to hardware differences, certain features and functionalities of KVM virtualization on IBM POWER8 differ from what is supported on AMD64 and Intel 64 systems.

For details on installing and using KVM virtualization on POWER8 systems, see the Virtualization Deployment and Administration Guide. (BZ#1531672)

NVIDIA GPU devices can now be used by multiple guests simultaneously

The NVIDIA vGPU feature is now supported on Red Hat Enterprise Linux 7. This enables dividing a vGPU-compatible NVIDIA GPU into multiple virtual devices referred to as **mediated devices**. By assigning mediated devices to guest virtual machines, these guests are able to share the performance of a single physical GPU.

To configure this feature, manually create a mediated device for the libvirt service to be able to use it as a vGPU. For details, see the Virtualization Deployment and Administration Guide. (BZ#1292451)

KASLR for KVM guests

Red Hat Enterprise Linux 7.5 introduces the Kernel Address Space Randomization (KASLR) feature for KVM guest virtual machines. KASLR enables randomizing the physical and virtual address at which the kernel image is decompressed, and thus prevents guest security exploits based on the location of kernel objects.

KASLR is activated by default, but can be deactivated on a specific guest by adding the **nokaslr** string to the guest's kernel command line.

Note that kernel crash dumps of guests with KASLR activated cannot be analyzed using the **crash** utility. To fix this, add the **<vmcoreinfo/>** element to the **<features>** section of the XML configuration files of your guests. However, KVM guests with **<vmcoreinfo/>** cannot be migrated to a

host system that does not support this element. This includes hosts that use Red Hat Enterprise Linux 7.4 and earlier (BZ#[1411490](#), BZ#[1395248](#))

Parallel decompression of ova files supported

With this release, the **pigz** and **pxz** decompression utilities are supported by the **virt-v2v** utility.

These utilities speed up extraction of **OVA** files compressed with the **gzip** and **xz** utilities on multi-processor machines. In addition, the command-line interfaces for **pigz** and **pxz** are fully compatible with the command-line interfaces for **gzip** and **xz**.

If **pigz** and **pxz** are installed, they are used by default. If **pigz** and **pxz** are not installed, there is no change to the extraction behavior. (BZ#[1448739](#))

SMAP now supported on Cannonlake guests

With this update, the Superior Mode Access Prevention (SMAP) feature is supported on guests that use the 7th Generation Intel Processors codenamed Cannonlake. This prevents malicious programs from forcing the kernel to use data from a user-space program, and thus increases the security of the guests.

To verify that your host CPU can provide SMAP for your guest, use the **virsh capabilities** command and look for the `<feature name='smap' />` string. (BZ#[1465223](#))

libvirt rebased to 3.9.0

The libvirt packages have been upgraded to version 3.9.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Sparse files are now preserved after moving them to or from another host.
- Response limits for remote procedure calls (RPCs) have been increased.
- Virtualized IBM POWER9 CPUs are now supported.
- Attaching devices to running guest virtual machines, also known as device hot plug, now supports more device types, such as input devices.
- The libvirt library has been secured against the CVE-2017-1000256 and CVE-2017-5715 security issues.
- VFIO-mediated devices now function more reliably. (BZ#[1472263](#))

virt-manager rebased to 1.4.3

The virt-manager packages have been upgraded to version 1.4.3, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- The virt-manager interface now displays the correct CPU models when creating a guest virtual machine that does not use the AMD64 and Intel 64 architectures.
- The default device selection has been optimized for guests using the IBM POWER, IBM z Systems, or the 64-bit ARM architectures.
- If an installed network card on the host system is compatible with single root I/O virtualization (SR-IOV), it is now possible to create a virtual network that lists a pool of available virtual functions of the selected SR-IOV-capable card.
- The selection of OS types and versions for a newly created guest has been expanded. (BZ#[1472271](#))

virt-what rebased to version 1.18

The virt-what packages have been updated to upstream version 1.18, which provides a number of bug fixes and enhancements over the previous version. Notably, the **virt-what** utility can now detect the following guest virtual machine types:

- Guests running on an 64-bit ARM host and booted using the Advanced Configuration and Power Interfaces.
- Guests running on the oVirt or Red Hat Virtualization hypervisor.
- Guests running on an IBM POWER7 host that uses logical partitioning (LPAR).
- Guests running on the FreeBSD bhyve hypervisor.
- Guests running on an IBM z Systems host that uses the KVM hypervisor.
- Guests emulated using the QEMU Tiny Code Generator (TCG).
- Guests running on the OpenBSD virtual machine monitor (VMM) service.
- Guests running on the Amazon Web Services (AWS) platform.
- Guests running on the Oracle VM Server for SPARC platform.

In addition, the following bugs have been fixed:

- The **virt-what** utility no longer fails on platforms that do not use the System Management BIOS (SMBIOS).
- **virt-what** now works correctly even if the \$PATH variable is not set. (BZ#[1476878](#))

tboot rebased to version 1.96

The tboot packages have been upgraded to upstream version 1.96, which fixes several bugs and adds various enhancements. Notable changes include:

- The OpenSSL library versions 1.1.0 and later are now supported for RSA key manipulation and ECDSA signature verification.
- Support has been added for event logs of Trusted Computing Group (TCG) trusted platform modules (TPMs).
- The x2APIC series of Advanced Programmable Interrupt Controller (APICs) is now supported.
- Additional checks have been added to prevent kernel images from being overwritten unintentionally.
- The **tboot** utility can no longer overwrite modules while moving them.
- A bug has been fixed that caused sealing and unsealing Amazon Simple Storage Service (S3) secrets to fail.
- Several null pointer dereference bugs have been fixed. (BZ#[1457529](#))

virt-v2v can convert VMware guests with snapshots

The **virt-v2v** utility has been enhanced to convert VMware guest virtual machines that have snapshots. Note that after the conversion, the status of such a guest is set to the top-most snapshot and the other snapshots are removed. (BZ#[1172425](#))

virt-rescue enhanced

This release of the **virt-rescue** utility includes the following enhancements:

- Ctrl+character sequences now act on commands run in **virt-rescue** and not on **virt-rescue** itself.
- The **-i** option allows users to mount all disks after inspecting the guest. (BZ#[1438710](#))

virt-v2v now converts Linux guests encrypted with LUKS

With this update, the **virt-v2v** utility can convert Linux guests installed with full-disk LUKS encryption, that is when all the partitions other than the **/boot** partition are encrypted.

Notes:

- The **virt-v2v** utility does not support conversion of Linux guests on partitions with other types of encryption schemes.
- The **virt-p2v** utility does not support conversion of Linux machines installed with full-disk LUKS encryption. (BZ#[1451665](#))

CAT support added to libvirt on specific CPU models

The **libvirt** service now supports Cache Allocation Technology (CAT) on specific CPU models. This enables guest virtual machines to have part of their host's CPU cache allocated for their vCPU threads.

For details on configuring this feature, see https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/virtualization_tuning_and_optimization_guide/index.html#sect_VTOG-vCPU_cache_reservation. (BZ#[1289368](#))

PTP device added to improve time synchronization of KVM guests

The PTP device has been added for KVM guest virtual machines. It enhances the **kvmclocks** service by preventing clock divergence between the host and the guest due to NTP adjustment. As a result, the PTP device ensures more reliable time synchronization between the KVM host and its guests.

For details on setting up the PTP device, see the Virtualization Deployment and Administration Guide. (BZ#[1379822](#))


CHAPTER 19. RED HAT ENTERPRISE LINUX 7.5 FOR ARM

Red Hat Enterprise Linux 7.5 for ARM introduces Red Hat Enterprise Linux 7.5 user space with an updated kernel, which is based on version 4.14 and is provided by the kernel-alt packages. The offering is distributed with other updated packages but most of the packages are standard Red Hat Enterprise Linux 7 Server RPMs. Installation ISO images are available on the [Customer Portal Downloads page](#).

For information about Red Hat Enterprise Linux 7.5 user space, see the [Red Hat Enterprise Linux 7 documentation](#). For information regarding the previous version, refer to [Red Hat Enterprise Linux 7.4 for ARM - Release Notes](#).

The following packages are provided as Development Preview in this release:

- libvirt (Optional channel)
- qemu-kvm-ma (Optional channel)



NOTE

KVM virtualization is a Development Preview on the 64-bit ARM architecture, and thus is not supported by Red Hat. For more information, see the [Virtualization Deployment and Administration Guide](#). Customers may contact Red Hat and describe their use case, which will be taken into consideration for a future release of Red Hat Enterprise Linux.

19.1. NEW FEATURES AND UPDATES

Core Kernel

- This update introduces the **qrwlock** queue write lock for 64-bit ARM systems. The implementation of this mechanism improves performance and prevents lock starvation by ensuring fair handling of multiple CPUs competing for the global task lock. This change also resolves a known issue, which was present in earlier releases and which caused soft lockups under heavy load.

Note that any kernel modules built for previous versions of Red Hat Enterprise Linux 7 for ARM (against the kernel-alt packages) must be rebuilt against the updated kernel. (BZ#1507568)

Security

USBGuard is now fully supported on 64-bit ARM systems

The **USBGuard** software framework provides system protection against intrusive USB devices by implementing basic whitelisting and blacklisting capabilities based on device attributes. Using **USBGuard** on 64-bit ARM systems, previously available as a Technology Preview, is now fully supported.

19.2. KERNEL CONFIGURATION CHANGES

HARDWARE ENABLEMENT

- Bluetooth (disabled)
- WIRELESS (disabled)
- CPU_IDLE (enabled)

- GPIO_DWAPB (enabled)
- I2C (enabled) - Designware, QUP, and XLP9XX
 - sensor support:
 - IIO drivers (disabled)
 - Accel sensors (disabled)
 - light + orientation + interrupt trigger (disabled)
- Input driver
 - mouse, synaptics, rmi4
- LED
 - Intel SS4200 (disabled)
- Generic IRQ CHIP (enabled)
- Hibernate (enabled)
- Clock Source DATA (enabled)
 - OSS_CORE (disabled)
 - all SND drivers (disabled)

Networking Driver Support

- Thunder2 driver (enabled)
- Amazon (enabled)
- Altera (disabled)
- ARC (disabled)
- Broadcom B44, BCMGENET, BNX2X_VLAN, CNIC (disabled)
- Hisilicon (enabled)
- cadence MACB (disabled)
- Chelsio T3 (disabled)
- Intel E1000 (disabled)
- Mellanox (enabled)
- myri10GE (disabled)
- Qlogic - qla2xxx, netxen_nic, Qed, Qede (enabled)
- Qualcomm - qcom_emac (enabled)

- Broadcom - bcm7xxx (disabled)

Infiniband Support

- CXBG4 (enabled)
- I40IW (enabled)
- MLX4 (enabled)
- MLX5 (enabled)
- IPOIB (enabled)
- IPOIB_CM (enabled)
- IPOIB_DEBUG (enabled)
- ISERT (enabled)
- SRP (enabled)
- SRPT (enabled)

CORE KERNEL SUPPORT

- Schedule Imbalance (enabled)
- 48 bit VA support (enabled)
- tick cpu accounting (disabled)
- Context Tracking (enabled)
- RCU NOCB (enabled)
- CGROUP-Hugetlb (enabled)
- CRIU (enabled)
- BPF_SYSCALL (disabled)
- PERF_USE_VMALLOC (disabled)
- HZ_100/HZ (enabled)
- NO_HZ_IDLE (disabled)
- NO_HZ_FULL (enabled)
- BPF_EVENTS (disabled)
- LZ4 compression (disabled)
- BTREE (enabled)
- CPUMASK_OFFSTACK (disabled)

- `DEBUG_INFO_DWARF4` (enabled)
- `SCHEDSTATS` (enabled)
- `Striaight DEVMEM` (disabled)
- `Transparent Hugepage (HTP)` (enabled)
- `ZSMaLLOC_STAT, IDLE_PAGE_TRACKING`(enabled)
- `PAGE_EXTENSION` and `PAGE_POISONING` (disabled)

Networking Stack Support

- `SLIP` - (enabled)
- `JME` (disabled)
- `IPVLAN` (disabled)
- `BPF_JIT` (disabled)
- `dccp` (disabled)
- `[ipv4] NET_FOU, Diag, CDG, NV` (disabled)
- `[ipv6] ILA` (disabled), `GRE` (enabled)
- `MAC80211` (disabled)
- `netfilter_conntrack` (enabled)

Graphic and GPU Support

- `DRM_I2C_SIL64` (disabled)
- `TTY`
 - `serial_nonstandard, cyclades, synclinkmp, synclink_gt, N_HDLC, serial_8250_MID` (enabled)
 - `fbdev` (enabled)
- `USB - PHY` (disabled)

Storage Support

- `Block scsi request` (enabled)
- `Block debugfs` (enabled)
- `Block Multi-Queue PCI` (enabled)
- `Block Multi-Queue VirtIO` (enabled)
- `Block Multi-Queue IOSched_deadline` (enabled)
- `MD Long Write` -(disabled)

- SCSI - ARCMSR, AM53C974, WD719x, BNX2X_FCOE, BNX2_ISCSI, ESAS2R (disabled)
- SCSI - HISI_SAS (enabled)
- SPI - QUP, SLP (enabled)
- SSB (disabled)

File Systems

- FS_DAX (enabled)
- BTRFS (disabled)
- Ceph (enabled)
- DLM (disabled)
- FSCAHE (disabled)
- GFS2 (disabled)
- Swap over NFS (disabled)
- NFS-FSCACHE (enabled)

Virtualization and KVM Support

- KVM_IRQCHIP, KVM_IRQ_ROUTING, KVM_MSI (enabled)
- Virtio - noiommu (enabled)

19.3. SUPPORT IN RED HAT SATELLITE

System management of Red Hat Enterprise Linux 7.5 for ARM is supported in Red Hat Satellite 6 but not in Red Hat Satellite 5.

19.4. KNOWN ISSUES

SELinux MLS policy is not supported with kernel version 4.14

SELinux Multi-Level Security (MLS) Policy denies unknown classes and permissions, and kernel version 4.14 in the kernel-alt packages recognizes the map permission, which is not defined in any policy. Consequently, every command on a system with active MLS policy and **SELinux** in enforcing mode terminates with the **Segmentation fault** error. A lot of **SELinux** denial warnings occurs on systems with active MLS policy and **SELinux** in permissive mode. The combination of **SELinux** MLS policy with kernel version 4.14 is not supported.

ipmitool communicates with BMC only when IPMI_SI=no on Cavium ThunderX

When starting **ipmi.service** with the **systemctl** command, the default configuration attempts to load the **ipmi_si** driver. On Cavium ThunderX systems, which do not have an IPMI SI device, **ipmi.service** incorrectly removes the **ipmi_devintf** driver. Consequently, the **ipmitool** utility in the kernel is not able to communicate with the Baseboard Management Controller (BMC). To work around this problem, edit the **/etc/sysconfig/ipmi** file and set the **IPMI_SI** variable as follows:

```
IPMI_SI=no
```

Then reboot the operating system if necessary. As a result, the correct drivers are loaded and **ipmitool** can communicate with BMC through the **/dev/ipmi0/** directory. (BZ#1448181)

Putting SATA ALPM devices into low power mode does not work correctly

When using the following commands to enable and disable low power mode for Serial Advanced Technology Attachment (SATA) devices using the Aggressive Link Power Management (ALPM) power management protocol on the 64-bit ARM systems, SATA does not work correctly:

```
tuned-adm profile powersave
```

```
tuned-adm profile performance
```

Consequently, SATA failures stop all disk I/O, and users have to reboot the operating system to fix it. To work around this problem, use one of the following options:

- Do not put the system into the powersave profile
- Check with your hardware vendor for firmware updates that might fix the bug with ALPM

(BZ#1430391)

Setting tuned to network-latency causes system hang on ARM

If the **tuned** profile is set to **network-latency** on the 64-bit ARM systems, the operating system becomes unresponsive, and the kernel prints a backtrace to the serial console. To work around this problem, do not set the **tuned** profile to **network-latency**. (BZ#1503121)

modprobe succeeds to load kernel modules with incorrect parameters

When attempting to load a kernel module with an incorrect parameter using the **modprobe** command, the incorrect parameter is ignored, and the module loads as expected on Red Hat Enterprise Linux 7 for ARM and for IBM Power LE (POWER9).

Note that this is a different behavior compared to Red Hat Enterprise Linux for traditional architectures, such as AMD64 and Intel 64, IBM z Systems and IBM Power Systems. On these systems, **modprobe** exits with an error, and the module with an incorrect parameter does not load in the described situation.

On all architectures, an error message is recorded in the **dmesg** output. (BZ#1449439)

19.5. BUG FIXES

The ld linker generates correct dynamic executables

Previously, the **ld** linker failed to create correct dynamic executables and terminated when invoked by the Go language compiler **go** on the 64-bit ARM architecture. The linker has been updated to correctly handle copy relocations. As a result, the linker no longer fails in the described situation. (BZ#1430743)

The ld linker generates correct dynamic relocations for constant data

Previously, the **ld** linker generated an incorrect kind of dynamic relocations for constant data shared between a library and executable on the 64-bit ARM architecture. As a consequence, the produced executable files wasted resources and terminated unexpectedly when the shared data was accessed. The linker has been updated to generate correct dynamic relocations, and the described problem no longer occurs. (BZ#1452170)

qrwlock is now enabled for 64-bit ARM systems

This update introduces the **qrwlock** queued read-write lock for 64-bit ARM systems. The

implementation of this mechanism improves performance and prevents lock starvation by ensuring fair handling of multiple CPUs competing for the global task lock. This change also resolves a known issue tracked in Red Hat Bugzilla #1454844, which was present in earlier releases and which caused soft lockups under heavy load.

Note that any kernel modules built for previous versions of Red Hat Enterprise Linux 7 for ARM (against the kernel-alt packages) must be rebuilt against the updated kernel.

CMA disabled by default

On 64-bit ARM Red Hat Enterprise Linux systems with memory limited to 1G or below, the Contiguous Memory Allocator (CMA) consumed large amount of memory, thus leaving insufficient memory for the rest of the kernel. Consequently, the out-of-memory (OOM) condition sometimes occurred in the kernel or certain user-space applications, such as Shared Memory in Linux (SHM)(/dev/shm).

The **CMA** support in the Red Hat Enterprise Linux kernel is now disabled by default for all architectures, and **CMA** no longer causes OOM.(BZ#[1519317](#))

CHAPTER 20. RED HAT ENTERPRISE LINUX 7.5 FOR IBM POWER LE (POWER9)

Red Hat Enterprise Linux 7.5 for IBM Power LE (POWER9) introduces Red Hat Enterprise Linux 7.5 user space with an updated kernel, which is based on version 4.14 and is provided by the kernel-alt packages. The offering is distributed with other updated packages but most of it is the standard Red Hat Enterprise Linux 7 Server RPMs. Installation ISO images are available on the [Customer Portal Downloads page](#).

For information about Red Hat Enterprise Linux 7.5 installation and user space, see the [Installation Guide](#) and other [Red Hat Enterprise Linux 7 documentation](#). For information regarding the previous version, refer to [Red Hat Enterprise Linux 7.4 for IBM Power LE \(POWER9\) - Release Notes](#).



NOTE

Bare metal installations on IBM Power LE using a USB drive require you to specify the **inst.stage2=** boot option manually at the boot menu. See the [Boot Options](#) chapter in the Installation Guide for detailed information.

20.1. NEW FEATURES AND UPDATES

Virtualization

- KVM virtualization is now supported on IBM POWER9 systems. However, due to hardware differences, certain features and functionalities differ from what is supported on AMD64 and Intel 64 systems. For details, see the [Virtualization Deployment and Administration Guide](#).

Platform Tools

- **OProfile** now includes support for the IBM POWER9 processor. Note that the **PM_RUN_INST_CMPL OProfile** performance monitoring event cannot be setup and should not be used in this version of **OProfile**. (BZ#1463290)
- This update adds support for the IBM POWER9 performance monitoring hardware events to **papi**. It includes basic PAPI presets for events, such as instructions (**PAPI_TOT_INS**) or processor cycles (**PAPI_TOT_CYC**). (BZ#1463291)
- This version of **libpfm** includes support for the IBM POWER9 performance monitoring hardware events. (BZ#1463292)
- **SystemTap** includes backported compatibility fixes necessary for the kernel.
- Previously, the **memcpy()** function from the GNU C Library (glibc) used unaligned vector load and store instructions on 64-bit IBM POWER systems. Consequently, when **memcpy()** was used to access device memory on POWER9 systems, performance would suffer. The **memcpy()** function has been enhanced to use aligned memory access instructions, to provide better performance for applications regardless of the memory involved on POWER9, without affecting the performance on previous generations of the POWER architecture. (BZ#1498925)

Security

USBGuard is now available as a Technology Preview on IBM Power LE (POWER9)

The **USBGuard** software framework provides system protection against intrusive USB devices by implementing basic whitelisting and blacklisting capabilities based on device attributes. **USBGuard** is now available as a Technology Preview on IBM Power LE (POWER9).

Note that USB is not supported on IBM z Systems, and the **USBGuard** framework cannot be provided on those systems.

20.2. KERNEL CONFIGURATION CHANGES

HARDWARE ENABLEMENT

- DEVFREQ_GOV_SIMPLE_ONDEMAND (enabled)
- GPIO IRQCHIP (enabled)
- HID plantronic (disabled)
- I2C sensors
 - JC42 (disabled)
 - NTC thermostat (enabled)
 - I2C MUX (enabled)

Networking Driver Support

- Broadcom B44 driver (disabled)
- Brocade BNA driver (disabled)
- Calxeda driver(disabled)
- IBM ethernet driver [ehea] (disabled)
- Intel E1000 driver (disabled)
- Mellanox driver [mlxsw] (disabled)
- Netronoma driver [NFP] (disabled)
- Qlogic [qla3xxx] driver (disabled)
- SFC falcon driver (disabled)
- Wireless (disabled)
 - WLAN (disabled)
 - Ath driver (disabled)
 - Ath10k driver (disabled)
 - Ath 9k driver (disabled)
 - Ath wil6210 (disabled)
 - Broadcom WLAN (disabled)
 - Broadcom brcm80211 (disabled)

- Intel WLAN (disabled)
- Intel iwlegacy (disabled)
- Intel iwlwifi (disabled)
- Marvell driver (disabled)
- Marvell mwiflex (disabled)
- Ralink WLAN driver (disabled)
- Ralink rt2x00 driver (disabled)
- Realtek driver (disabled)
- Realtek rt1818x driver (disabled)
- Realtek rtwifi driver (disabled)
- NVME driver + target driver (enabled)
- ptp 1588 driver (disabled)
- s390 HMC driver (disabled)
- RTL8192e driver (disabled)
- RTL8712u driver (disabled)
- Serial UARTLITE driver (enabled)
- USB LED trigger USBPORT (disabled)
- USBIP driver (disabled)
- Power Mgt Deubg + Adv Debug + Sleep Debug (enabled)

CORE KERNEL SUPPORT

- Sched Imbalance patchset (enabled)
- OPTprobes, kprobe on ftrace (enabled)
- 64bit Aligned Access (disabled)
- Arch Soft Dirty (enabled)
- Arch MMAP Rnd Compat (disabled)
- SWIOTLB (disabled)
- Crypto: akcipher, rsa (enabled)
- Compression:
 - Kernel gzip support (enabled)

- Kernel XZ support (enabled)
- Locking: Mutex spin on owner (enabled in debugging kernel)
- Function Tracer (enabled)
- Dynamic Ftrace (enabled)
- Ftrace mcount record (enabled)
- Common kernel Libraries
 - Rational (enabled)
 - Btree (enabled)
 - libfdt (enabled)
 - parman (disabled)
- MM
 - NO_BOOTMEM (enabled)
 - MOVABLE NODE (enabled)
 - HMM (Hetrogenous Memory Management) (enabled)
 - HMM Mirrored (enabled)
 - Coherent Device Memory (CDM) (enabled)
 - Zone Device (enabled)
- IMA (enabled)
- YAMA (disabled)

Networking Stack Support

- Compact Netlink Msg (disabled)
- BPF_JIT (enabled)
- DCCP (disabled)
- CCIDS (disabled)
- IPv6 NF target NPT (disabled)
- Mac80211 (disabled)

Desktop, Graphic, and GPU Support

- DRM_DP_AUX_CHARDEV (enabled)
- STK1160 video usb driver (disabled)

- V412 BUF2_DMA_SG (enabled)

Storage Support

- DAX (disabled)
- NVDIMM + PFN + DAX (enabled)
- SCSI
 - 3Ware 9xxx driver (disabled)
 - 3Ware sAS driver (disabled)
 - ARCMSR driver (disabled)
 - AIC79xx driver (disabled)
 - Broadcom Bnx2x driver (enabled)
 - Broadcom Bnx2 driver (disabled)
 - QED driver (disabled)
 - QEDI driver (disabled)

File Systems

- BTRFS (disabled)
- DLM (disabled)
- GFS2 DLM locking support (disabled)

Virtualization and KVM Support

- vhost [vsock] (disabled)
- VMWare vsock (disabled)

20.3. SUPPORT IN RED HAT SATELLITE

System management of Red Hat Enterprise Linux 7.5 for IBM POWER LE (POWER9) is supported in Red Hat Satellite 6 but not in Red Hat Satellite 5.

20.4. KNOWN ISSUES

SELinux MLS policy is not supported with kernel version 4.14

SELinux Multi-Level Security (MLS) Policy denies unknown classes and permissions, and kernel version 4.14 in the kernel-alt packages recognizes the map permission, which is not defined in any policy. Consequently, every command on a system with active MLS policy and **SELinux** in enforcing mode terminates with the **Segmentation fault** error. A lot of **SELinux** denial warnings occurs on systems with active MLS policy and **SELinux** in permissive mode. The combination of **SELinux** MLS policy with kernel version 4.14 is not supported.

kdump saves the vmcore only if mpt3sas is blacklisted

When **kdump** kernel loads the **mpt3sas** driver, the **kdump** kernel crashes and fails to save the **vmcore** on certain POWER9 systems. To work around this problem, blacklist **mpt3sas** from the **kdump** kernel environment by appending the **module_blacklist=mpt3sas** string to the **KDUMP_COMMANDLINE_APPEND** variable in the **/etc/sysconfig/kdump** file:

```
KDUMP_COMMANDLINE_APPEND="irqpoll maxcpus=1 ... module_blacklist=mpt3sas"
```

Then restart the **kdump** service to pick up the changes to the configuration file by running the **systemctl restart** command as the **root** user:

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

As a result, **kdump** is now able to save the **vmcore** on the POWER9 systems. (BZ#1496273)

Recovering from OOM situation fails due to incorrect function of OOM-killer

Recovering from an out-of-memory (OOM) situation does not work correctly on systems with large amounts of memory. Kernel's OOM-killer kills the process using the most memory and frees the memory to be used again. However, sometimes the OOM-killer does not wait long enough before killing a second process. Eventually, the OOM-killer kills all the processes on the system and logs this error:

```
Kernel panic - not syncing: Out of memory and no killable processes...
```

If this happens, the operating system must be rebooted. There is no available workaround. (BZ#1405748)

HTM is disabled for guests running on IBM POWER systems

The Hardware Transactional Memory (HTM) feature currently prevents migrating guest virtual machines from IBM POWER8 to IBM POWER9 hosts, and has therefore been disabled by default. As a consequence, guest virtual machines running on IBM POWER8 and IBM POWER9 hosts cannot use HTM, unless the feature is manually enabled.

To do so, change the default **pseries-rhel7.5** machine type of these guests to **pseries-rhel7.4**. Note that guests configured this way cannot be migrated from an IBM POWER8 host to an IBM POWER9 host. (BZ#1525599)

Migrating guests with huge pages from IBM POWER8 to IBM POWER9 fails

IBM POWER8 hosts can only use 16MB and 16GB huge pages, but these huge-page sizes are not supported on IBM POWER9. As a consequence, migrating a guest from an IBM POWER8 host to an IBM POWER9 host fails if the guest is configured with static huge pages.

To work around this problem, disable huge pages on the guest and reboot it prior to migration. (BZ#1538959)

modprobe succeeds to load kernel modules with incorrect parameters

When attempting to load a kernel module with an incorrect parameter using the **modprobe** command, the incorrect parameter is ignored, and the module loads as expected on Red Hat Enterprise Linux 7 for ARM and for IBM Power LE (POWER9).

Note that this is a different behavior compared to Red Hat Enterprise Linux for traditional architectures, such as AMD64 and Intel 64, IBM z Systems and IBM Power Systems. On these systems, **modprobe** exits with an error, and the module with an incorrect parameter does not load in the described situation.

On all architectures, an error message is recorded in the **dmesg** output. (BZ#1449439)

20.5. BUG FIXES

kdump no longer hangs due to the attempts to read the memory from on-board devices

On the little-endian variants of IBM Power Systems hardware, the **kdump** mechanism became unresponsive because the kernel attempted to read the memory from on-board devices such as the GPU, and include it as a part of the **vmcore**. This update fixes **kexec-tools** to skip the on-board devices when attempting to read the memory during **kdump**. As a result, **kdump** now works correctly, the **vmcore** is saved to disk and the operating system reboots as expected. (BZ#1478049)

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 Systems, 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 7.5 that have a significant impact on users.

CHAPTER 23. GENERAL UPDATES

runc notifies systemd about user-specified CPU quota limits

Previously, the **runc** program did not notify **systemd** about user-specified CPU quota limits when a container was started. Consequently, **systemd** was unaware of the user-specified limits, and therefore the CPU quota was reset to the default value (unlimited) during the **systemctl daemon-reload** operation. With this update, **runc** now notifies **systemd** about user-specified CPU quota limits when a container is started, and the described problem no longer occurs. (BZ#1455071)

Segmentation faults in applications because of only non-existent paths in LD_LIBRARY_PATH no longer happen

Previously, when the **LD_LIBRARY_PATH** environment variable contained only non-existent paths, the dynamic loader produced a segmentation fault. Consequently, applications terminated unexpectedly with a segmentation fault at startup in the described situation. The dynamic loader has been fixed. As a result, applications no longer terminate unexpectedly in the described situation.

Note that updating the **glibc** package is enough to fix this bug for any affected applications. (BZ#1443236)

The setup package now creates the tape group with the correct group number

Previously, when installing the **setup** package, the **tape** group was created with an ID that was inconsistent with all other versions of Red Hat Enterprise Linux. With this update, the group ID has been changed from **30** to the standard **33**. As a result, fresh installations of the operating system now have the correct group number for the **tape** group.

On previously installed systems affected by this problem:

1. Edit the group ID in the **/etc/group** and **/etc/gshadow** files.
2. Change the group ownership for all files owned by the former **tape** group. (BZ#1433020)

CHAPTER 24. AUTHENTICATION AND INTEROPERABILITY

The IdM LDAP server no longer becomes unresponsive when resolving an AD user takes a long time

When the System Security Services Daemon (SSSD) took a long time to resolve a user from a trusted Active Directory (AD) domain on the Identity Management (IdM) server, the IdM LDAP server sometimes exhausted its own worker threads. Consequently, the IdM LDAP server was unable to respond to further requests from SSSD clients or other LDAP clients. This update adds a new API to SSSD on the IdM server, which enables identity requests to time out. Also, the IdM LDAP extended identity operations plug-in and the Schema Compatibility plug-in now support this API to enable canceling requests that take too long. As a result, the IdM LDAP server can recover from the described situation and keep responding to further requests. (BZ#[1415162](#), BZ#[1473571](#), BZ#[1473577](#))

Application configuration snippets in `/etc/krb5.conf.d/` are now automatically read in existing configurations

Previously, Kerberos did not automatically add support for the `/etc/krb5.conf.d/` directory to existing configurations. Consequently, application configuration snippets in `/etc/krb5.conf.d/` were not read unless the user added the include statement for the directory manually. This update modifies existing configurations to include the appropriate `includedir` line pointing to `/etc/krb5.conf.d/`. As a result, applications can rely on their configuration snippets in `/etc/krb5.conf.d`.

Note that if you manually remove the `includedir` line after this update, successive updates will not add it again. (BZ#[1431198](#))

`pam_mkhomedir` can now create home directories under `/`

Previously, the `pam_mkhomedir` module was unable to create subdirectories under the `/` directory. Consequently, when a user with a home directory in a non-existent directory under `/` attempted to log in, the attempt failed with this error:

```
Unable to create and initialize directory '/<directory_path>'.
```

This update fixes the described problem, and `pam_mkhomedir` is now able to create home directories in this situation.

Note that even after applying this update, SELinux might still prevent `pam_mkhomedir` from creating the home directory, which is the expected SELinux behavior. To ensure `pam_mkhomedir` is allowed to create the home directory, modify the SELinux policy using a custom SELinux module, which enables the required paths to be created with the correct SELinux context. (BZ#[1509338](#))

Kerberos operations depending on KVNO in the keytab file no longer fail when a RODC is used

The `adcli` utility did not handle the key version number (KVNO) properly when updating Kerberos keys on a read-only domain controller (RODC). Consequently, some operations, such as validating a Kerberos ticket, failed because no key with a matching KVNO was found in the keytab file. With this update, `adcli` detects if a RODC is used and handles the KVNO accordingly. As a result, the keytab file contains the right KVNO, and all Kerberos operations depending on this behavior work as expected. (BZ#[1471021](#))

`krb5` properly displays errors about PKINIT misconfiguration in single-realm KDC environments

Previously, when Public Key Cryptography for Initial Authentication in Kerberos (PKINIT) was misconfigured, the `krb5` package did not report the incorrect configuration to the administrator. For example, this problem occurred when the deprecated `pkinit_kdc_ocsp` option was specified in the

`/etc/krb5.conf` file. With this update, krb5 exposes PKINIT initialization failures when only one realm is specified in the Kerberos key distribution center (KDC). As a result, single-realm KDCs report PKINIT misconfiguration properly. (BZ#1460089)

Certificate System no longer incorrectly logs `ROLE_ASSUME` audit events

Previously, Certificate System incorrectly generated the `ROLE_ASSUME` audit event for certain operations even if no privileged access occurred for a user. Consequently, the event was incorrectly logged. The problem has been fixed and `ROLE_ASSUME` events are no longer logged in the mentioned scenario. (BZ#1461524)

Updated attributes in `CERT_STATUS_CHANGE_REQUEST_PROCESSED` audit log event

Previously, the `CERT_STATUS_CHANGE_REQUEST_PROCESSED` audit event in log files contained the following attributes:

- **ReqID** - The requester ID
- **SubjectID** - The subject ID of the certificate

For consistency with other audit events, the attributes have been modified and now contain the following information:

- **ReqID** - The request ID
- **SubjectID** - The requester ID (BZ#1461217)

Signed audit log verification now works correctly

Previously, due to improper logging system initialization and incorrect signature calculation by the verification tool, signed audit log verification could fail on the first log entry and after log rotation. With this update, the logging system and the verification tool have been fixed. As a result, signed audit log verification now works correctly in the mentioned scenarios. (BZ#1404794)

Certificate System now validates the banner file

A previous version of Certificate System introduced a configurable access banner - a custom message to be displayed in the PKI console at the start of every secure session. The contents of this banner were not validated, which could cause a `JAXBUnmarshalException` error if the message contained invalid UTF-8 characters. With this update, the contents of the banner file are validated both on server startup and on client requests. If the file is found to contain invalid UTF-8 characters on server startup, the server will not start. If invalid characters are found when a client requests the banner, the server will return an error message and not send the banner to the client. (BZ#1446579)

The TPS subsystem no longer fails when performing a symmetric key changeover on a HSM

Previously, attempting to perform a symmetric key changeover with the master key on a Hardware Security Module (HSM) token failed with an error reported by the Certificate System Token Processing System (TPS) subsystem. This update fixes the way the master key on a HSM is used to calculate the new key set, allowing the TPS to successfully upgrade a token key set when the master resides on a HSM. The fix is currently verified with the G&D SmartCafe 6.0 HSM. (BZ#1465142)

Certificate System CAs no longer display an error when handing subject DN's without a CN component

Previously, an incoming request missing the Common Name (CN) component caused a `NullPointerException` on the Certificate Authority (CA) because the implementation expected the CN to be present in the subject Distinguished Name (DN) of the Certificate Management over CMS (CMC). This update allows the CA to handle subject DN without a CN component, preventing the exception from being thrown. (BZ#1474658)

The pki-server-upgrade utility no longer fails if target files are missing

A bug in the **pki-server-upgrade** utility caused it to attempt to locate a non-existent file. As a consequence, the upgrade process failed to complete, and could possibly leave the PKI deployment in an invalid state. With this update, **pki-server-upgrade** has been modified to correctly handle cases where target files are missing, and PKI upgrades now work correctly. (BZ#[1479663](#))

The Certificate System CA key replication now works correctly

A previous update to one of the key unwrapping functions introduced a requirement for a key usage parameter which was not being supplied at the call site, which caused lightweight Certificate Authority (CA) key replication to fail. This bug has been fixed by modifying the call site so that it supplies the key usage parameter, and lightweight CA key replication now works as expected. (BZ#[1484359](#))

Certificate System no longer fails to import PKCS #12 files

An earlier change to PKCS #12 password encoding in the Network Security Services (NSS) caused Certificate System to fail to import PKCS #12 files. As a consequence, the Certificate Authority (CA) clone installation could not be completed. With this update, PKI will retry a failed PKCS #12 decryption with a different password encoding, which allows it to import PKCS #12 files produced by both old and new versions of NSS, and CA clone installation succeeds. (BZ#[1486225](#))

The TPS user interface now displays the token type and origin fields

Previously, the **tps-cert-find** and **tps-cert-show** Token Processing System (TPS) user interface utilities did not display the token type and origin fields which were present in the legacy TPS interface. The interface has been updated and now displays the missing information. (BZ#[1491052](#))

Certificate System issued certificates with an expiration date later than the expiration date of the CA certificate

Previously, when signing a certificate for an external Certificate Authority (CA), Certificate System used the **ValidityConstraint** plug-in. Consequently, it was possible to issue certificates with a later expiry date than the expiry date of the issuing CA. This update adds the **CAValidityConstraint** plug-in to the registry so that it becomes available for the enrollment profiles. In addition, the **ValidityConstraint** plug-in in the **caCMCcaCert** profile has been replaced with the **CAValidityConstraint** plug-in which effectively sets the restrictions. As a result, issuing certificates with an expiry date later than the issuing CA is no longer allowed. (BZ#[1518096](#))

CA certificates without SKI extension no longer causes issuance failures

A previous update of Certificate System incorrectly removed a fallback procedure, which generated the Issuer Key Identifier. Consequently, the Certificate Authority (CA) failed to issue certificates if the CA signing certificate does not have the Subject Key Identifier (SKI) extension set. With this update, the missing procedure has been added again. As a result, issuing certificates no longer fails if the CA signing certificate does not contain the SKI extension. (BZ#[1499054](#))

Certificate System correctly logs the user name in CMC request audit events

Previously, when Certificate System received a Certificate Management over CMS (CMC) request, the server logged an audit event with the **SubjectID** field set to **\$NonRoleUser\$**. As a result, administrators could not verify who issued the request. This update fixes the problem, and Certificate System now correctly logs the user name in the mentioned scenario. (BZ#[1506819](#))

The Directory Server trivial word check password policy now works as expected

Previously, when you set a **userPassword** attribute to exactly the same value as an attribute restricted by the **passwordTokenMin** setting with the same length, Directory Server incorrectly allowed the password update operation. With this update, the trivial word check password policy feature now correctly verifies the entire user attribute value as a whole, and the described problem no longer occurs. (BZ#[1517788](#))

The `pkidestroy` utility now fully removes instances that are started by the `pki-tomcatd-nuxwdog` service

Previously, the `pkidestroy` utility did not remove Certificate System instances that used the `pki-tomcatd-nuxwdog` service as a starting mechanism. As a consequence, administrators had to migrate `pki-tomcatd-nuxwdog` to the service without watchdog before using `pkidestroy` to fully remove an instance. The utility has been updated, and instances are correctly removed in the mentioned scenario.

Note that if you manually removed the password file before running `pkidestroy`, the utility will ask for the password to update the security domain. (BZ#1498957)

The Certificate System deployment archive file no longer contains passwords in plain text

Previously, when you created a new Certificate System instance by passing a configuration file with a password in the `[DEFAULT]` section to the `pkispawn` utility, the password was visible in the archived deployment file. Although this file has world readable permissions, it is contained within a directory that is only accessible by the Certificate Server instance user, which is `pkiuser`, by default. With this update, permissions on this file have been restricted to the Certificate Server instance user, and `pkispawn` now masks the password in the archived deployment file.

To restrict access to the password on an existing installation, manually remove the password from the `/etc/sysconfig/pki/tomcat/<instance_name>/<subsystem>/deployment.cfg` file, and set the file's permissions to `600`. (BZ#1532759)

ACIs with the `targetfilter` keyword work correctly

Previously, if an Access Control Instruction (ACI) in Directory Server used the `targetfilter` keyword, searches containing the `geteffective` rights control returned before the code was executed for template entries. Consequently, the `GetEffectiveRights()` function could not determine the permissions when creating entries and returned false-negative results when verifying an ACI. With this update, Directory Server creates a template entry based on the provided `geteffective` attribute and verifies access to this template entry. As a result, ACIs in the mentioned scenario work correctly. (BZ#1459946)

Directory Server searches with a scope set to `one` have been fixed

Due to a bug in Directory Server, searches with a scope set to `one` returned all child entries instead of only the ones that matched the filter. This update fixes the problem. As a result, searches with scope `one` only return entries which are matching the filter. (BZ#1511462)

Clear error message when sending TLS data to a non-LDAPS port

Previously, Directory Server decoded TLS protocol handshakes sent to a port that was configured to use plain text as an `LDAPMessage` data type. However, decoding failed and the server reported the misleading **BER was 3 bytes, but actually was <greater>** error. With this update, Directory Server detects if TLS data is sent to a port configured for plain text and returns the following error message to the client:

```
Incoming BER Element may be misformed. This may indicate an attempt to use
TLS on a plaintext port, IE ldaps://localhost:389. Check your client
LDAP_URI settings.
```

As a result, the new error message indicates that an incorrect client configuration causes the problem. (BZ#1445188)

Directory Server no longer logs an error if not running the `cleanallruv` task

After removing a replica server from an existing replication topology without running the `cleanallruv`

task, Directory Server previously logged an error about not being able to replace referral entries. This update adds a check for duplicate referrals and removes them. As a result, the error is no longer logged. (BZ#[1434335](#))

Using a large number of CoS templates no longer slow down the virtual attribute processing time

Due to a bug, using a large number of Class of Service (CoS) templates in Directory Server increased the virtual attribute processing time. This update improves the structure of the CoS storage. As a result, using a large number of CoS templates no longer increases the virtual attribute processing time. (BZ#[1523183](#))

Directory Server now handles binds during an online initialization correctly

During an online initialization from one Directory Server master to another, the master receiving the changes is temporarily set into a referral mode. While in this mode, the server only returns referrals. Previously, Directory Server incorrectly generated these bind referrals. As a consequence, the server could terminate unexpectedly in the mentioned scenario. With this update, the server correctly generates bind referrals. As a result, the server now correctly handles binds during an online initialization. (BZ#[1483681](#))

The `dirsrv@.service` meta target is now linked to `multi-user.target`

Previously, the `dirsrv@.service` meta target had the `Wants` parameter set to `dirsrv.target` in its systemd file. When you enabled `dirsrv@.service`, this correctly enabled the service to the `dirsrv.target`, but `dirsrv.target` was not enabled. Consequently, Directory Server did not start when the system booted. With this update, the `dirsrv@.service` meta target is now linked to `multi-user.target`. As a result, when you enable `dirsrv@.service`, Directory Server starts automatically when the system boots. (BZ#[1476207](#))

The `memberof` plug-in now logs all update attempts of the `memberof` attribute

In certain situations, Directory Server fails to update the `memberof` attribute of a user entry. In this case, the `memberof` plug-in logs an error message and forces the update. In the previous Directory Server version, the second try was not logged if it was successful. Consequently, the log entries were misleading, because only the failed attempt was logged. With this update, the `memberof` plug-in also logs the successful update if the first try failed. As a result, the plug-in now logs the initial failure, and the subsequent successful retry as well. (BZ#[1533571](#))

The Directory Server password policies now work correctly

Previously, subtree and user password policies did not use the same default values as the global password policy. As a consequence, Directory Server incorrectly skipped certain syntax checks. This bug has been fixed. As a result, the password policy features work the same for the global configuration and the subtree and user policies. (BZ#[1465600](#))

A buffer overflow has been fixed in Directory Server

Previously, if you configured an attribute to be indexed and imported an entry that contained a large binary value into this attribute, the server terminated unexpectedly due to a buffer overflow. The buffer has been fixed. As a result, the server works as expected in the mentioned scenario. (BZ#[1498980](#))

Directory Server now sends the password expired control during grace logins

Previously, Directory Server did not send the expired password control when an expired password had grace logins left. Consequently, clients could not tell the user that the password was expired or how many grace logins were left. The problem has been fixed. As a result, clients can now tell the user if a password is expired and how many grace logins remain. (BZ#[1464505](#))

An unnecessary global lock has been removed from Directory Server

Previously, when the `memberOf` plug-in was enabled and users and groups were stored in separate back ends, a deadlock could occur. An unnecessary global lock has been removed and, as a result, the deadlock no longer occurs in the mentioned scenario. (BZ#1501058)

Replication now works correctly with TLS client authentication and FIPS mode enabled

Previously, if you used TLS client authentication in a Directory Server replication environment with Federal Information Processing Standard (FIPS) mode enabled, the internal Network Security Services (NSS) database token differed from a token on a system with FIPS mode disabled. As a consequence, replication failed. The problem has been fixed, and as a result, replication agreements with TLS client authentication now work correctly if FIPS mode is enabled. (BZ#1464463)

Directory Server now correctly sets whether virtual attributes are operational

The `pwdpolycysubentry` subtree password policy attribute in Directory Server is flagged as operational. However, in the previous version of Directory Server, this flag was incorrectly applied to following virtual attributes that were processed. As a consequence, the search results were not visible to the client. With this update, the server now resets the attribute before processing the next virtual attribute and Class of Service (CoS). As a result, the expected virtual attributes and CoS are now returned to the client. (BZ#1453155)

Backup now succeeds if replication was enabled and a changelog file existed

Previously, if replication was enabled and a changelog file existed, performing a backup on this master server failed. This update sets the internal options for correctly copying a file. As a result, creating a backup now succeeds in the mentioned scenario. (BZ#1476322)

Certificate System updates the revocation reason correctly

Previously, if a user temporarily lost a smart card token, the administrator of a Certificate System Token Processing System (TPS) in some cases changed the status of the certificate from **on hold** to **permanently lost** or **damaged**. However, the new revocation reason did not get reflected on the CA. With this update, it is possible to change a certificate status from **on hold** directly to **revoked**. As a result, the revocation reason is updated correctly. (BZ#1500474)

A race condition has been fixed in the Certificate System clone installation process

In certain situations, a race condition arose between the LDAP replication of security domain session objects and the execution of an authenticated operation against a clone other than the clone where the login occurred. As a consequence, cloning a Certificate System installation failed. With this update, the clone installation process now waits for the security domain login to finish before it enables the security domain session objects to be replicated to other clones. As a result, the clone installation no longer fails. (BZ#1402280)

Certificate System now uses strong ciphers by default

With this update, the list of enabled ciphers has been changed. By default, only strong ciphers, which are compliant with the Federal Information Processing Standard (FIPS), are enabled in Certificate System.

RSA ciphers enabled by default:

- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256

- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA

Note that the **TLS_RSA_WITH_AES_128_CBC_SHA** and **TLS_RSA_WITH_AES_256_CBC_SHA** ciphers need to be enabled to enable the **pkispawn** utility to connect to the LDAP server during the installation and configuration.

ECC ciphers enabled by default:

- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256

In addition, the default ranges of the **sslVersionRangeStream** and **sslVersionRangeDatagram** parameters in the `/var/lib/pki/<instance_name>/conf/server.xml` file now use only TLS 1.1 and TLS 1.2 ciphers. (BZ#[1539125](#))

The pkispawn utility no longer displays incorrect errors

Previously, during a successful installation of Certificate System, the **pkispawn** utility incorrectly displayed errors related to deleting temporary certificates. The problem has been fixed, and the error messages no longer display if the installation succeeds. (BZ#[1520277](#))

The Certificate System profile configuration update method now correctly handles backslashes

Previously, a parser in Certificate System removed backslash characters from the configuration when a user updated a profile. As a consequence, affected profile configurations could not be correctly imported, and issuing certificates failed or the system issued incorrect certificates. Certificate System now uses a parser that handles backslashes correctly. As a result, profile configuration updates import the configuration correctly. (BZ#[1541853](#))

CHAPTER 25. CLUSTERING

Pacemaker correctly implements fencing and unfencing for Pacemaker remote nodes

Previously, Pacemaker did not implement unfencing for Pacemaker remote nodes. As a consequence, Pacemaker remote nodes remained fenced even if a fence device required unfencing. With this update, Pacemaker correctly implements both fencing and unfencing for Pacemaker remote nodes, and the described problem no longer occurs. (BZ#1394418)

Pacemaker now probes guest nodes

Important update for users of guest nodes.

Pacemaker now probes guest nodes, which are Pacemaker remote nodes created using the **remote-node** parameter of a resource such as **VirtualDomain**. If users were previously relying on the fact that probes were not done, the probes may fail, potentially causing fencing of the guest node. If a guest node cannot run a probe of a resource (for example, if the software is not even installed on the guest), then the location constraint banning the resource from the guest node should have the **resource-discovery** option set to **never**, the same as would be required with a cluster node or remote node in the same situation. (BZ#1489728)

The **pcs resource cleanup** command no longer generates unnecessary cluster load

The **pcs resource cleanup** command cleans the records of failed resource operations that have been resolved. Previously, the command probed all resources on all nodes, generating an unnecessary load on cluster operation. With this fix, the command probes only the resources for which a resource operation failed. The previous functionality of the **pcs resource cleanup** command has been replaced by the new **pcs resource refresh** command, which probes all resources on all nodes. For information on cluster resource cleanup, see https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/high_availability_add-on_reference/#s1-resource_cleanup-HAAR. (BZ#1508351)

Warning generated when user specifies **action** attribute for stonith device

Previously, it was possible for a user to set an action attribute for a stonith device, even though this option is deprecated and is not recommended as it can cause unexpected fencing. The following fixes have been implemented:

- When a user tries to set an **action** option of a stonith device with the CLI, this generates a warning message along with the instructions to use the **--force** flag to set this attribute.
- The **pcsd** Web UI now displays a warning message next to **action** option field.
- The output of the **pcs status** command displays a warning when a stonith device has the **action** option set. (BZ#1421702)

It is now possible to enable stonith agent debugging without specifying the **--force** flag

Previously, attempting to enable debugging of a stonith agent by setting the **debug** or **verbose** parameters required that the user specify the **--force** flag. With this fix, using the **--force** flag is no longer necessary. (BZ#1432283)

The **fence_ilo3** resource agent no longer has a default value of **cycle** for the **action** parameter

Previously, the **fence_ilo3** resource agent had a default value of **cycle** for the **action** parameter.

This value is unsupported, as it may cause data corruption. The default value for this parameter is now **onoff**. Additionally, a warning is now displayed in the output of the **pcs status** command and the web UI if a stonith device has its **method** option set to **cycle**. (BZ#[1519370](#), BZ#[1523378](#))

Pacemaker no longer starts up when `sbd` is enabled but not started successfully by `systemd`

Previously, if **sbd** did not start properly, **systemd** would still start Pacemaker. This would lead to **sbd** poison pill triggered reboots not being performed without this being detected by **fence_sbd** and, in the case of quorum-based watchdog fencing, the nodes losing quorum would not self-fence either. With this fix, if **sbd** does not come up properly Pacemaker is not started. This should prevent all sources of data corruption due to **sbd** not coming up. (BZ#[1525981](#))

A fenced node in an ‘sbd’ setup now shuts down reliably

Previously, when a node received an ‘off’ via the poison pill mechanism used by ‘sbd’ on a shared disk, the node would be likely to reboot instead of powering off. With this fix, receiving an ‘off’ will power off the node. Receiving a ‘reset’ will reboot the node. If the node is not able to perform the software-driven reboot or power off properly, the watchdog is going to trigger and the action performed is what the watchdog device is configured to. A fenced node in an ‘sbd’ setup now shuts down reliably if the watchdog device is configured to power off the node, and fencing is requesting ‘off’ via the poison pill mechanism on a shared disk. (BZ#[1468580](#))

IPaddr2 resource agent now finds NIC for IPv6 addresses with 128 netmask

Previously, the **IPaddr2** resource agent failed to find the NIC for IPv6 addresses with 128 netmask. This fix corrects that issue. (BZ#[1445628](#))

portblock agent no longer yields excessive unnecessary messages

Previously, the **portblock** agent would flood the `/var/log/messages` file with monitoring messages that provided no useful information. With this fix, the `/var/log/messages` file contains more limited logging output from the **portblock** agent. (BZ#[1457382](#))

/var/run/resource-agents directory now persists across reboots

Previously, the `/var/run/resource-agents` directory, created at installation of the **resource-agents** package, was not persistent across reboots. With this fix, the directory is now present after a reboot. (BZ#[1462802](#))

CHAPTER 26. COMPILER AND TOOLS

Package selection now works in `system-config-kickstart`

A bug in the `system-config-kickstart` graphical Kickstart file creation utility caused the package selection to be unavailable because the tool could not download package information from repositories. This bug is now fixed, and you can now configure package selection in `system-config-kickstart` again. (BZ#[1272068](#))

NVMe devices no longer show up as `Unknown` in `parted` and `Anaconda`

Previously, any Non-Volatile Memory Express (NVMe) devices were not being recognized by the `Anaconda` installer and the `parted` storage configuration tool during the installation, and were instead being labeled as `Model: Unknown (unknown)`. This update backports an upstream patch that enables recognition of these devices, and they are now being correctly identified as `NVMe Device (nvme)` during installation. (BZ#[1316239](#))

`DBD:MySQL` now sends and receives smaller integers correctly on big-endian platforms

Previously, the `DBD:MySQL` Perl driver incorrectly handled integers smaller than 64 bits on big endian platforms. Consequently, tests for prepared statements failed for certain variable sizes on the IBM z Systems architecture. This bug has been fixed, and the described problem no longer occurs. (BZ#[1311646](#))

The `version` Perl module now supports tainted input and tainted version objects

Previously, the `version` module of Perl was unable to correctly parse tainted input. Consequently, when building a version object from a tainted variable, the `version->new()` method reported the `Invalid version format (non-numeric data)` error. This update adds support for parsing tainted input and for printing tainted version objects and strings. (BZ#[1378885](#))

The `HTTP:Daemon` Perl module now supports IPv6

Previously, the `HTTP:Daemon` Perl module did not support IPv6 addresses. Consequently, when running an `HTTP:Daemon:SSL` server on an IPv6 address, the server terminated unexpectedly on an attempt to print the IPv6 address with an `Arg length for inet_ntoa` error message. With this update, the `HTTP:Daemon` module has been ported from the `IO:Socket:INET` to the `IO:Socket:IP` module. As a result, `HTTP:Daemon` handles IPv6 addresses as expected. (BZ#[1413065](#))

`GDB` shows inline function names in breakpoint listing

Previously, the `GDB` debugger showed caller function names instead of inlined callee function names when listing breakpoints. As a consequence, `GDB` users were not able to identify breakpoints placed on inline functions from the function name. `GDB` has been extended to store names of inline callee functions when breakpoints are placed. As a result, `GDB` now correctly displays names of inline functions when listing breakpoints. (BZ#[1228556](#))

Relocation failures at module load time due to wrong `gcc` alignment fixed

Previously, `GCC` generated code containing `.toc` sections with 2^0 alignment. As a consequence, relocation failures could occur at module load time. `GCC` has been changed to generate `.toc` sections aligned to 2^3 . This fix eliminates most cases of occurrence of this bug. (BZ#[1487434](#))

The `istream::sentry` object from the `gcc C++` standard library no longer throws exceptions

Previously, the `istream::sentry` object from the `gcc C++` standard library did not properly handle exceptions that happen while skipping whitespace. As a consequence, an unexpected exception could

occur in the object's code. The constructor for the **sentry** class has been fixed to catch the exceptions and update the error state of the **istream** object appropriately. (BZ#1469384)

Multiple fixes in **gdb** on IBM Power

Previously, various features of the **gdb** debugger have been broken on the IBM Power architecture:

- Record and replay functionality was not available and resulted in error messages or not restoring the previous register values.
- Printing short vector return values resulted in wrong values displayed.
- Single stepping over atomic sequences failed to actually step over them - the program counter did not change.

This update fixes these features. (BZ#[1480498](#), BZ#1480496, BZ#1480497)

GDB no longer crashes when dumping core from a process that terminates

Previously, the **GDB** debugger did not consider that a process can be terminated while **GDB** is dumping it into a core file. As a consequence, when a dumped program terminated after receiving an unexpected **SIGKILL** signal, the **gcore** utility terminated unexpectedly as well. With this update, **GDB** has been extended to handle this situation. As a result, **GDB** and the **gcore** command no longer terminate unexpectedly and create invalid core files. (BZ#1493675)

GDB can again dump memory protected by the **VM_DONTDUMP** flag

Previous changes to the GNU Debugger **GDB** made the behavior of the **gcore** command more similar to the behavior of the Linux kernel when dumping process memory to increase data security. Consequently, users of **GDB** could not dump memory protected by the **VM_DONTDUMP** flag. The new **set dump-excluded-mappings** setting has been added to **GDB** to enable dumping of memory with this flag. As a result, users can dump the whole process memory with **GDB** again. (BZ#[1518243](#))

Programs using the **CLONE_PTRACE** flag on threads now run under **strace**

Previously, programs which set the **CLONE_PTRACE** flag on new threads caused undefined behavior of the **strace** tool, because it uses the **ptrace()** function for its operation. As a consequence, such programs could be neither traced nor executed properly. The **strace** tool has been modified to ignore threads with an unexpected **CLONE_PTRACE** flag. As a result, programs which use **CLONE_PTRACE** execute properly under **strace**. (BZ#1466535)

exiv2 rebased to version 0.26

The **exiv2** packages have been upgraded to upstream version 0.26, which provides a number of bug fixes and enhancements over the previous version. Notably, **exiv2** now contains:

- CMake support for Visual Studio
- Recursive File Dump
- ICC Profile Support
- The **exiv2** command for metadata piping
- Lens File for user lens definitions
- User defined lens types
- WebP Support

For the complete changelog, see <http://www.exiv2.org/changelog.html#v0.26>. (BZ#1420227)

gssproxy fixed to properly update ccaches

Previously, the gssproxy package did not correctly handle the key version number (kvno) incrementation in Kerberos credential caches (ccaches). As a consequence, stale ccaches were not properly overwritten. This update fixes these problems in gssproxy ccache caching. As a result, ccaches are now properly updated, and the caching prevents excessive requests for updates. (BZ#1488629)

gcc on the little-endian variant of IBM Power Systems architecture no longer creates unused stack frames

Previously, using the `-pg -mprofile=kernel` options of the **gcc** compiler on the little-endian variant of IBM Power Systems architecture could result in unused stack frames being generated for leaf functions. The **gcc** compiler has been fixed and the unused stack frames no longer occur in this situation. (BZ#1468546)

Several bugs fixed in gssproxy

This update fixes several bugs in the gssproxy package. The bug fixes include preventing potential memory leaks and concurrency problems. (BZ#1462974)

The BFD library regains the ability to convert binary addresses to source code positions

A previous enhancement to the **BFD** library from the binutils package caused a bug in parsing the DWARF debug information. As a consequence, **BFD** and all tools using it, such as **gprof** and **perf**, were unable to convert binary file addresses to positions in source code. With this update, **BFD** has been modified to prevent the described problem. As a result, **BFD** can now convert addresses in binary files into positions in source code as expected.

Note that tools that use the **BFD** library must be relinked in order to take advantage of this fix. (BZ#1465318)

Applications using vector registers for passing arguments work again

Previously, the dynamic loader in the GNU C library (**glibc**) contained an optimization which avoided saving and restoring vector registers for 64-bit Intel and AMD architectures. Consequently, applications compiled for these architectures and using unsupported vector registers for passing function arguments, not adhering to the published x86-64 psABI specification, could fail and produce unexpected results. This update changes the dynamic loader to use the **XSAVE** and **XSAVEC** context switch CPU instructions, preserving more CPU state, including all vector registers. As a result, applications using vector registers for argument passing, in ways which are not supported by the x86-64 psABI specification, work again. (BZ#1504969)

curl now properly resets the HTTP authentication state

Prior to this update, the authentication state was not reset properly when an HTTP transfer finished or when the `'curl_easy_reset()'` function was called. Consequently, the **curl** tool did not send the request body to the following URL. With this update, the authentication state is reset properly when an HTTP transfer is done or when `curl_easy_reset()` is called, and the described problem no longer occurs. (BZ#1511523)

The strip utility works again

Previously, the **BFD** library missed a NULL pointer check on the IBM z Systems architecture. As a consequence, running the **strip** utility caused a segmentation fault. This bug has been fixed, and **strip** now works as expected. (BZ#1488889)

Importing python modules generated by f2py now works properly

Previously, when dynamic linking loader was configured to load symbols globally, a segmentation fault

occurred when importing any python module generated by the **f2py** utility. This update renames the **PyArray_API** symbol to **_numpy_f2py_ARRAY_API**, which prevents potential conflicts with the same symbol in the multiarray module. As a result, importing modules generated by **f2py** no longer leads to a segmentation fault. (BZ#[1167156](#))

mailx is not encoding multi-byte subjects properly

Previously, the **mailx** mail user agent did not split non-ASCII message headers on multi-byte character boundaries when encoding into the Multipurpose Internet Mail Extension (MIME) standard. As a consequence, the headers were incorrectly decoded. This update modifies the MIME encoding function so that it splits headers into encoded words on multi-byte character boundaries. As a result, **mailx** now sends messages with headers that can be properly decoded. (BZ#1474130)

The --all-logs option now works as expected in sosreport

Previously, the **--all-logs** option was ignored by the **apache**, **nscd**, and **logs** plug-ins of the **sosreport** utility. This bug has been fixed, and the mentioned plug-ins now correctly handle **--all-logs**. Note that when using **--all-logs**, it is impossible to limit the size of the log with the **--log-size** option, which is an expected behavior. (BZ#1183243)

Python scripts can now correctly connect to HTTPS servers through a proxy, while explicitly setting the port

The Python standard library provided in Red Hat Enterprise Linux was previously updated to enable certificate verification by default. However, a bug prevented Python scripts using the standard library from connecting to HTTPS servers using a proxy when explicitly setting the port to connect to. The same bug also prevented users from using the bootstrap script for registration with Red Hat Satellite 6 through a proxy. This bug is now fixed, and scripts can now connect to HTTPS servers and register using Red Hat Satellite as expected. (BZ#[1483438](#))

CHAPTER 27. DESKTOP

Stylus of Dell Canvas 27 fixed

Previously, Dell Canvas 27 contained a Wacom tablet in which the ranges were offset by default. As a consequence, the stylus mapped to the upper left quarter of the screen. Red Hat Enterprise Linux 7.5 supports the stylus of the Dell Canvas 27, making sure coordinates are accurately reported. As a result, the cursor is placed directly under the tip of the stylus as required. (BZ#1507821)

llvmpipe crashes on IBM Power Systems

On the little-endian variant of IBM Power Systems architecture, a race-condition in **GNOME Shell** code previously caused that, the LLVM engine for Mesa, **llvm-private**, terminated unexpectedly. This update disables threading in the JavaScript engine which prevents the segmentation fault from occurring. As a result, **llvm-private** no longer crashes on IBM Power Systems. (BZ#1523121)

CHAPTER 28. FILE SYSTEMS

NFS shares no longer become unresponsive after a TCP connection is closed

Previously, NFS clients sometimes entered a 60 second **TIME_WAIT** period after initiating the TCP disconnect sequence. This happened only when TCP timestamps were disabled on the connection. During the waiting period, the client was unable to reconnect the NFS TCP connection.

Due to waiting in the **TIME_WAIT** period, the NFS mount points were unresponsive, an **rpciod** kernel thread was using 100% CPU, and the **retrans** number in the output of the **nfsstat -r** command was becoming a very large number. In addition, NFS mounts with lower values of the **timeo** and **retrans** options could cause I/O errors.

With this update, the NFS TCP connection is able to reconnect immediately after a disconnect sequence using a different source port. As a result, NFS mounts no longer become unresponsive and **rpciod** no longer causes a high system load after a connection is closed. (BZ#1479043)

CHAPTER 29. HARDWARE ENABLEMENT

genwqe-tools updated for IBM Power Systems ppc64 and ppc64le architectures

The genwqe-tools packages have been updated for IBM Power Systems and the little-endian variant of IBM Power Systems. This enhancement update includes the following backported fixes from genwqe-tools master branch:

- the **adler32** detect corruption checksum now returns correction on the **deflateSetDictionary()** function
- the **deflateSetDictionary()** function now returns error on NULL dictionary as required by the spec file
- The debugger has been removed from the **zpipe_rnd.c** file
- Potential overflow in expression has been avoided
- Out of bounds access and possible resource leak have been fixed
- To simplify contributions, a Contributor License Agreement (CLA) has been changed to the Developer's Certificate of Origin (DCO)
- Potential security hole has been resolved
- The Failure of the Hardware Accelerator Tool genwqe_cksum which causes EEH, has been resolved

Users of genwqe-tools are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. (BZ#1456492)

Hardware utility tools now correctly identify recently released hardware

Prior to this update, obsolete ID files caused that recently released hardware connected to a computer was reported as unknown. To fix this bug, PCI, USB, and vendor device identification files have been updated. As a result, hardware utility tools now correctly identify recently released hardware. (BZ#1489281)

CHAPTER 30. INSTALLATION AND BOOTING

The installer no longer crashes when you select an incomplete IMSM RAID array during manual partitioning

Previously, if the system being installed had a storage drive which was previously part of an Intel Matrix (IMSI) RAID array which was broken at the time of the installation, the disk was displayed as **Unknown** in the **Installation Destination** screen in the graphical installer. If you attempted to select this drive as an installation target, the installer crashed with the **An unknown error has occurred** message. This update adds proper handling for such drives, and allows you to use them as standard installation targets. (BZ#1465944)

Installer now accepts additional time zone definitions in Kickstart files

Starting with Red Hat Enterprise Linux 7.0, **Anaconda** switched to a different, more restrictive method of validating time zone selections. This caused some time zone definitions, such as **Japan**, to be no longer valid despite being acceptable in previous versions, and legacy Kickstart files with these definitions had to be updated or they would default to the **Americas/New_York** time zone.

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

Note that this change only applies to the `timezone` Kickstart command. The time zone selection in the graphical and text-based interactive interfaces remains unchanged. Existing Kickstart files for Red Hat Enterprise Linux 7 that had valid time zone selections do not require any updates. (BZ#1452873)

Proxy configuration set up using a boot option now works correctly in Anaconda

Previously, proxy configuration made in the boot menu command line using the `proxy=` option was not correctly applied when probing remote package repositories. This was caused by an attempt to avoid a refresh of the **Installation Source** screen if network settings were changed. This update improves the installer logic so that proxy configuration now applies at all times but still avoids blocking the user interface on settings changes. (BZ#1478970)

FIPS mode now supports loading files over HTTPS during installation

Previously, installation images did not support FIPS mode (`fips=1`) during installation where a Kickstart file is being loaded from an HTTPS source (`inst.ks=https://<location>/ks.cfg`). This release implements support for this previously missing functionality, and loading files over HTTPS in FIPS mode works as expected. (BZ#1341280)

Network scripts now correctly update `/etc/resolv.conf`

Network scripts have been enhanced to update the `/etc/resolv.conf` file correctly. Notably:

- The scripts now update the `nameserver` and `search` entries in the `/etc/resolv.conf` file after the `DNS*` and `DOMAIN` options, respectively, have been updated in the `ifcfg-*` files in the `/etc/sysconfig/network-scripts/` directory
- The scripts now also update the order of `nameserver` entries after it has been updated in the `ifcfg-*` files in `/etc/sysconfig/network-scripts/`
- Support for the `DNS3` option has been added
- The scripts now correctly process duplicate and randomly omitted `DNS*` options (BZ#1364895)

Files with the `.old` extension are now ignored by network scripts

Network scripts in Red Hat Enterprise Linux contain a regular expression which causes them to ignore **ifcfg-*** configuration files with certain extensions, such as **.bak**, **.rpmnew** or **.rpmold**. However, the **.old** extension was missing from this set, despite being used in documentation and in common practice. This update adds the **.old** extension into the list, which ensures that script files which use it will be ignored by network scripts as expected. (BZ#[1455419](#))

Bridge devices no longer fail to obtain an IP address

Previously, bridge devices sometimes failed to obtain an IP address from the DHCP server immediately after system startup. This was caused by a race condition where the **ifup-eth** script did not wait for the Spanning Tree Protocol (STP) to complete its startup. This bug has been fixed by adding a delay that causes **ifup-eth** to wait long enough for STP to finish starting. (BZ#[1380496](#))

The `rhel-dmesg` service can now be disabled correctly

Previously, even if the **rhel-dmesg.service** was explicitly disabled using **systemd**, it continued to run anyway. This bug has been fixed, and the service can now be disabled correctly. (BZ#[1395391](#))

CHAPTER 31. KERNEL

kdump can now capture a vmcore with nokaslr set

When using **nokaslr** and **crashkernel=xxM,high** options, a bug in the implementation of **nokaslr** prevented the **kdump** mechanism from capturing a **vmcore** file. This fix ensures that if **nokaslr** is set, the original loading address of the kernel is returned. As a result, **kdump** can now collect a **vmcore** when Kernel Address Space Layout Randomization (KASLR) is compiled in the kernel, but disabled with **nokaslr**, and high memory is specified in the **crashkernel** parameter. (BZ#1467561)

MPOL_PREFERRED policy now works with Transparent Huge Pages (THP) with optimal performance

Allocating memory on node 1 with the **MPOL_PREFERRED** policy did not work with Transparent Huge Pages (THP) enabled, but always fell back to the node 0 local node. Consequently, workload performance for multinode systems was significantly impacted. The backported patch ensures **MPOL_PREFERRED** policy with non-local node is respected, and system performance is back to optimal. (BZ#1476709)

A cgroups deadlock has been fixed

In certain circumstances when using **cgroups**, a system deadlock occurred due to a race condition. This update adds a work queue that fixes the race condition, which prevents the deadlock from happening. (BZ#1476040)

System no longer becomes unresponsive when DM thin provisioning is used on top of a loop device

Previously, system sometimes became unresponsive when Device Mapper (DM) thin provisioning was used on top of a loop device. With this update, memory allocation now uses correct gfp mask. As a result, the described problem no longer occurs. (BZ#1469247)

KASLR now no longer causes mirroring of kernel memory to non-mirrored regions

Prior to this update, with specified mirrored memory regions and kernel address space layout randomization (KASLR) enabled kernel memory could be located into non-mirrored memory regions. As a consequence, non-mirrored memory regions became unmovable. With this update, Kernel memory location is restricted from mirror regions. As a result, KASLR no longer causes mirroring of kernel memory to non-mirrored regions. (BZ#1446684)

Users now receive message with prompt to remove white space characters in the /etc/kdump.conf

Previously, one or more leading white space characters before a **kdump** configuration item in the **/etc/kdump.conf** caused incorrect **kdump** configuration. With this update, an error message with prompt to remove the leading white space characters return to users, and **kdump** no longer fails due to the described behavior. (BZ#1476219)

An application with large .bss segment on IBM POWER Systems will no longer cause random segmentation faults

Previously, on IBM POWER Systems architectures, an application with large **.bss** segment could cause the dynamic linker to terminate unexpectedly. As a consequence, an application launched with the dynamic linker could randomly cause segmentation faults. With this update, the **ELF_ET_DYN_BASE** value has been increased to 4GB for 64-bit implementations and 4MB for 32-bit implementations on this architecture. As a result, an application with large **.bss** segment on IBM POWER Systems architectures will not lead to random segmentation faults. (BZ#1432288)

Kernel no longer consumes excessive amounts of resources to calculate load

Previously, the kernel calculated load for every task group, including empty task groups, which consumed an excessive amount of system resources on systems with a large number of processes. This update prevents the kernel from calculating the load of empty task groups, which reduces the system load in the described circumstances. (BZ#1460641)

Cpuset is now able to restore the effective CPU mask after a pair of offline and online events

Prior to this update, the **cpuset** filesystem, which confines processes to processor and memory node subsets, had one bitmap set enabled for CPUs used in a cpuset. As a consequence, a CPU offline event followed by a CPU online event caused the affected CPU to be removed from all non-root cpusets. With this update, cpuset has two bitmap sets enabled. As a result, cpuset is now able to properly track CPU online or offline events to restore the effective CPU mask as long as the **-o cpuset_v2_mode** mount flag is used when mounting cpuset cgroup. (BZ#947004)

Access to /proc/[pid]/maps is now significantly faster

Previously, the time to locate a task of a stack Virtual Memory Area (VMA) in the **[stack:TID]** annotation scaled directly with the number of active tasks in the system. As a consequence, the more tasks were running in the system, the slower it was to correctly annotate the stack VMA, which causes slowed access to the **/proc/[pid]/maps** files. With this update, the **[stack:TID]** annotation is no longer used. As a result, access to **/proc/[pid]/maps** is now significantly faster, particularly when a lot of tasks is running in the system. (BZ#1448534)

fadump no longer fails to restart

Previously, fadump stopped during DLPAR memory remove operation and then started to restart. Under certain circumstances fadump failed to restart. With this update, the described problem no longer occurs. (BZ#1438695)

makedumpfile can now map page table entries correctly

On some virtual machines running on HP hardware, it was impossible to correctly obtain the physical address of the virtual machine's memory, causing the **makedumpfile** utility to fail with an error similar to:

```
readmem: Can't convert a virtual address(ffffffffb21158a0) to physical address
```

The problem happened because **file_size** was incorrectly calculated, preventing the **readpage_elf()** function from working properly. This update fixes the calculation of **file_size** on these systems, ensuring that a **vmcore** file can be collected, and the **makedumpfile --mem-usage** command estimates the **vmcore** size correctly. (BZ#1448861)

Asymmetric groups are used for overlapping scheduling domains

Previously, scheduling group construction on certain Non-Uniform Memory Access (NUMA) systems negatively influenced thread migration. This situation adversely affected the performance when a task could not be migrated to a neighboring NUMA node. With this update, asymmetric groups are used for overlapping scheduling domains to solve the problem. (BZ#1373534)

The KASLR no longer causes kernel to become unresponsive while booting the system

Previously, the kernel sometimes became unresponsive on certain SGI UV systems when the Kernel Address Space Layout Randomization (KASLR) feature was enabled. As a consequence, the systems were unable to boot. With this update, the kernel does not attempt to adapt the size of the direct mapping when KASLR is enabled. As a result, the system now boots normally and the described problem no longer occurs. (BZ#1457046)

Unplugging a Wacom tablet with ExpressKeys no longer causes the operating system to reboot

When some Wacom tablets were unplugged from a running GNOME session on Red Hat Enterprise Linux 7.4, the operating system rebooted within five seconds. This problem was initially observed on Wacom model 27QHD devices. This update ensures that the tablet can be unplugged without causing the operating system to reboot. (BZ#1462363)

Setting `memory.kmem.limit_in_bytes` no longer causes a problem when removing that memory cgroup later

Previously, setting the cgroup `memory.kmem.limit_in_bytes` parameter caused a problem when that memory cgroup was later removed. The problem occurred when an attempt was made to merge the memory cgroup `kmem` cache, which was not handled properly. This update disables `kmem` cache merging for memory cgroups by backporting the current upstream code, which no longer uses this functionality. (BZ#1442618)

The `sha1-avx2` encryption algorithm is now re-enabled

Due to a read-beyond error (when the code attempts to read memory outside of its boundary), the `sha1-avx2` encryption algorithm was disabled. With this update, the problem has been resolved, and administrators may now use `sha1-avx2`. (BZ#1469200)

VXLAN rebased to version 4.14

The VXLAN driver has been upgraded to upstream version 4.14, which provides a number of bug fixes over the previous version. Notable changes include the following:

- The tunnel source IP address is used in route lookups.
- VXLAN Generic Protocol Extension (VXLAN-GPE) now uses the correct Internet Assigned Numbers Authority (IANA) for User Datagram Protocol (UDP) port.
- The `VNI 0xffffffff` value can now be used.
- A race condition on tunnel removal has been fixed.
- Static forwarding database (fdb) entries now behave consistently with Linux bridge. (BZ#1467280)

CHAPTER 32. NETWORKING

Network operation persists when `ip6mr` unregisters an already unregistered device

Previously, the **IPv6 multicast routing (`ip6mr`)** code tried to unregister an already unregistered device. As a consequence, a bug was reported in the **syslog** causing the network operation to stop. With this update, **ip6mr** no longer unregisters devices that are already marked as unregistered. As a result, no more bugs are reported in **syslog**, and the network operation persists in the described scenario. (BZ#1445046)

Sending big files through **VTI** no longer fails

Previously, when sending a big file through **Virtual Tunnel Interface (VTI)** failed because **VTI** did not handle **Path Maximum Transmission Unit (PMTU)**. As a consequence, files with greater size than the **PMTU** size could not be sent. This update adds **PMTU** handling. As a result, **PMTU** can be updated in Tx path, and the described problem no longer occurs. (BZ#1467521)

L2TP with **IPv6** encapsulation now works in name space

Previously, using **Layer 2 Tunneling Protocol (L2TP)** with **IPv6** encapsulation did not support name space. As a consequence, **L2TP** could not be used in name space. With this update, **L2TP** with **IPv6** encapsulation is now aware of name space, and the described problem no longer occurs. (BZ#1465711)

Flushing **ARP** entries no longer fails

Previously, trying to flush an incomplete or failed **Address Resolution Protocol (ARP)** entry had no effect. As a consequence, the incomplete **ARP** entry remained there, and in some cases caused problems for debugging systems or networks. This update allows for the removal of an incomplete or failed **ARP** entry. As a result, users can now get an **ARP** table as expected. (BZ#1383691, BZ#1469945)

Using `cls_matchall` with classful queue disciplines no longer causes the kernel to crash

Previously, the matchall classifier (`cls_matchall`) did not assign the **classic** option to a packet. As a consequence, the kernel terminated unexpectedly when trying to use `cls_matchall` with classful queueing disciplines (**classful qdiscs**), such as Hierarchical Token Bucket (HTB) or Class Based Queueing (CBQ). With this update, when `cls_matchall` processes **classid**, **classid** is assigned to a packet. As a result, `cls_matchall` with **classful qdiscs** can now be used successfully and the user-provided value of **classid** is no longer ignored in the described scenario.

For more details on the kernel actions related to **classid**, see the **OPTIONS** section in the `tc-matchall (8)` man page. (BZ#1460213)

ICMP error packets are no longer lost when a user connects to a closed **SCTP** port

Previously, when trying to connect to a closed Stream Control Transmission Protocol (SCTP) port, an **Internet Control Message Protocol (ICMP)** error reply from the server was lost. This occurred only with **Network Interface Cards (NICs)** that used non-linear buffers to receive data. As a consequence, for a connection to a closed SCTP port, the user was waiting until a timeout instead of getting the **connection refused** error message from the server immediately. With this update, the received data is handled in a linear way and the **ICMP** error reply is not lost. As a result, the user receives the corresponding **ICMP** error in the described situation. (BZ#1450529)

SCTP now selects the right source address

Previously, when using a secondary IPv6 address, Stream Control Transmission Protocol (SCTP) selected the source address based on the best prefix matching with the destination address. As a consequence, in some cases, a packet was sent through an interface with the wrong IPv6 address. With

this update, SCTP uses the address that already exists in the routing table for this specific route. As a result, SCTP uses the expected IPv6 address as the source address when secondary addresses are used on a host. (BZ#1460106)

Device reference held by iptables CLUSTERIP target is now properly released on namespace deletion

Previously, the **iptables CLUSTERIP** target held a direct reference to the network device specified as input device in the associated rule. When that rule inside a namespace was deleted, the corresponding reference was not released. As a consequence, upon namespace deletion, dangling references held by the **CLUSTERIP** target sometimes prevented deletion of network devices contained in the namespace. For this reason, it was not possible to create a device with the same name and the related memory was not freed. With this update, the **CLUSTERIP** target rule reference does not hold the related device but its index. As a result, when deleting a namespace, all the rules and references related to this namespace are also cleared properly. (BZ#1472892)

The nftables configuration files are no longer publicly readable

Previously, during installation in the **RPM** file, the **nftables** configuration file mode bits were not adjusted accordingly. As a consequence, the configuration templates in the **/etc/nftables** directory and the **etc/sysconfig/nftables.conf** main configuration file were publicly readable. With this update, the file mode bits are explicitly set to correct values when installing the configuration files. As a result, the user can now install the configuration files with the correct permissions.

Note that the configuration files which are not modified by the administrator, are replaced with configuration files with the correct permissions.

The modified configuration files are not replaced. In that case, for **/etc/sysconfig/nftables.conf**, an rpmnew file is created which has the correct permissions. For any files in **/etc/nftables**, no rpmnew file is created, and the user must manually set the permissions. (BZ#1451404)

The Ready to read events are now correctly sent to an application when SENDER_DRY_EVENTS is enabled

Previously, when enabling the **SENDER_DRY_EVENTS** notifications or when the Stream Control Transmission Protocol (SCTP) Partial Reliability triggered the removal of a chunk, the SCTP stack flagged an event that it was already generated and sent it to an application. However, the flag was not removed afterwards. As a consequence, the application missed the **ready to read** event. With this update, the stack does not flag the event in such cases anymore. As a result, the **ready to read** events are now correctly dispatched to an application. (BZ#1442784)

SCTP statistics now available

Previously, the stream control transmission protocol (SCTP) statistics parser could not handle the **/proc/net/sctp/snmp** source file. As a consequence, users were not able to see the statistic information. Parsing of the SCTP statistics has been fixed. As a result, the SCTP statistics are now available to users. (BZ#1329338)

The firewalld service daemon no longer hangs in the rmmmod process

Previously, some network device drivers, specifically some **wi-fi** and **IP over InfiniBand Network Interface Cards (IPoIB NICs)** drivers, held **contrack** entries associated with untracked packets for an unlimited amount of time. As a consequence, at removal time, the **contrack** kernel module was in a busy loop waiting for these entries to be freed. This led to the **rmmmod nf_contrack** module consuming 100% of the CPU usage causing **firewalld** to hang at shutdown time. With this update, the new kernel removes support for the **notrack** **contrack** entries, and **contrack** no longer waits for such entries to be freed. As a result, the **firewalld** shutdown no longer hangs. (BZ#1317099)

CHAPTER 33. SECURITY

When `firewalld` starts, `net.netfilter.nf_conntrack_max` is no longer reset to default if its configuration exists

Previously, `firewalld` reset the `nf_conntrack` settings to their default values when it was started or restarted. As a consequence, the `net.netfilter.nf_conntrack_max` setting was restored to its default value. With this update, each time `firewalld` starts, it reloads `nf_conntrack` sysctls as they are configured in `/etc/sysctl.conf` and `/etc/sysctl.d`. As a result, `net.netfilter.nf_conntrack_max` maintains the user-configured value. (BZ#1462977)

Tomcat can now be started using `tomcat-jsvc` with SELinux in enforcing mode

In Red Hat Enterprise Linux 7.4, the `tomcat_t` unconfined domain was not correctly defined in the SELinux policy. Consequently, the Tomcat server cannot be started by the `tomcat-jsvc` service with SELinux in enforcing mode. This update allows the `tomcat_t` domain to use the `dac_override`, `setuid`, and `kill` capability rules. As a result, Tomcat is now able to start through `tomcat-jsvc` with SELinux in enforcing mode. (BZ#1470735)

SELinux now allows `vdsm` to communicate with `lldpad`

Prior to this update, SELinux in enforcing mode denied the `vdsm` daemon to access `lldpad` information. Consequently, `vdsm` was not able to work correctly. With this update, a rule to allow a `virttd_t` domain to send data to a `lldpad_t` domain through the `dgram` socket has been added to the selinux-policy packages. As a result, `vdsm` labeled as `virttd_t` can now communicate with `lldpad` labeled as `lldpad_t` if SELinux is set to enforcing mode. (BZ#1472722)

OpenSSH servers without Privilege Separation no longer crash

Prior to this update, a pointer had been dereferenced before its validity was checked. Consequently, OpenSSH servers with the `Privilege Separation` option turned off crashed during the session cleanup. With this update, pointers are checked properly, and OpenSSH servers no longer crash while running without `Privilege Separation` due the described bug.

Note that disabling `OpenSSH Privilege Separation` is not recommended. (BZ#1488083)

The `clevis luks bind` command no longer fails with the DISA STIG-compliant password policy

Previously, passwords generated as part of the `clevis luks bind` command were not compliant with the Defense Information Systems Agency Security Technical Implementation Guide (DISA STIG) password policy set in the `pwquality.conf` file. Consequently, `clevis luks bind` failed on DISA STIG-compliant systems in certain cases. With this update, passwords are generated using a utility designed to generate random passwords that pass the password policy, and `clevis luks bind` now succeeds in the described scenario. (BZ#1500975)

WinSCP 5.10 now works properly with OpenSSH

Previously, OpenSSH incorrectly recognized WinSCP version 5.10 as older version 5.1. As a consequence, the compatibility bits for WinSCP version 5.1 were enabled for WinSCP 5.10, and the newer version did not work properly with OpenSSH. With this update, the version selectors have been fixed, and WinSCP 5.10 now works properly with OpenSSH servers. (BZ#1496808)

SFTP no longer allows to create zero-length files in read-only mode

Prior to this update, the `process_open` function in the OpenSSH SFTP server did not properly prevent write operations in read-only mode. Consequently, attackers were allowed to create zero-length files. With this update, the function has been fixed, and the SFTP server no longer allows any file creation in read-only mode. (BZ#1517226)

CHAPTER 34. SERVERS AND SERVICES

Internal buffer locks no longer cause deadlocks in `libdb`

Previously, the `libdb` database did not lock its internal buffers in the correct order when it accessed pages located in an off-page duplicate (OPD) tree while processing operations on a cursor. A writer process accessed first the primary tree and then the OPD tree while a reader process did the same in the reverse order. When a writer process accessed a page from the primary tree while a reader process accessed a page from the OPD tree, the processes were unable to access the page from the other tree because the pages were simultaneously locked by the other process. This consequently caused a deadlock in `libdb` because neither of the processes released their locks. With this update, the `btree` version of the `cursor->get` method has been modified to lock the tree's pages in the same order as the writing methods, that is, the primary tree first and the OPD tree second. As a result, deadlocks in `libdb` no longer occur in the described scenario. (BZ#[1349779](#))

Weekly log rotations are now triggered more predictably

Weekly log rotations were previously performed by the `logrotate` utility when exactly 7 days (604800 seconds) elapsed since the last rotation. Consequently, if the `logrotate` command was triggered by a cron job slightly sooner, the rotation was delayed until the next run. With this update, weekly log rotations ignore the exact time. As a result, when the day counter advances by 7 or more days since the last rotation, a new rotation is triggered. (BZ#[1465720](#))

`ghostscript` no longer crashes while processing large PDF files

Previously, processing large PDF files could cause the `ghostscript` utility to terminate unexpectedly under certain rare circumstances. With this update, an internal `ghostscript` virtual machine limit, `DEFAULT_VM_THRESHOLD`, has been increased, and the described problem no longer occurs. In addition, processing of large files is now slightly faster. (BZ#[1479852](#))

Converting large PDF files to PNG with `ghostscript` no longer fails

Due to a bug in the upstream source code, converting large PDF files to the PNG format using the `ghostscript` utility failed under certain rare circumstances. This bug has been fixed, and the described problem no longer occurs. (BZ#[1473337](#))

`krfb` no longer crashes when unable to bind to an IPv6 port

Previously, connecting to the `krfb` application with a VNC client when `krfb` could not bind to an IPv6 port, `krfb` terminated unexpectedly. This update fixes the improper handling of uninitialized IPv6 socket, and applications built on the `libvncserver` library now deal with the unsuccessful attempt to listen on an IPv6 port correctly. (BZ#[1314814](#))

`mod_nss` properly detects the threading model in Apache to improve performance

Previously, the `mod_nss` module was not detecting the threading model properly in Apache. Consequently, users experienced slower performance because the TLS Session ID was not maintained across handshakes and a new session ID was generated for each handshake. This update fixes the threading model detection. As a result, TLS Session IDs are now properly cached, which eliminates the described performance problems. (BZ#[1461580](#))

`atd` no longer runs with 100% CPU utilization nor fills system log

Previously, the `atd` daemon of the `at` utility handled incorrectly some types of broken jobs, particularly jobs of non-existent users. As a consequence, `atd` used up all available CPU resources and filled the system log by messages sent with unlimited frequency. With this update, the handling of the broken jobs by `atd` has been fixed and the problem does not occur anymore. (BZ#[1481355](#))

ReaR now provides a more helpful error message when `grub2-efi-x64-modules` is missing

Previously, an attempt to create a **ReaR** backup on UEFI systems using the **rear mkrescue** and **rear mkbackup** commands failed due to a missing grub2-efi-x64-modules package, which is not installed by default but is required by **ReaR** to generate a GRUB image. The commands failed with the following error message:

```
ERROR: Error occurred during grub2-mkimage of BOOTX64.efi
```

This message proved to be confusing and unhelpful. With this update, the error will still appear in the same circumstances, but it will point out how to fix the problem:

```
WARNING: /usr/lib/grub/x86_64-efi/moddep.lst not found, grub2-mkimage will
likely fail. Please install the grub2-efi-x64-modules package to fix this.
```

As the updated message explains, you must install the missing grub2-efi-x64-modules package before you can create a **ReaR** backup on a system with UEFI firmware. (BZ#[1492177](#))

ReaR no longer fails to determine disk size during a mkrescue operation

Previously, the **ReaR** (Relax-and-Recover) utility sometimes encountered a failure while querying partition sizes when saving the disk layout due to a race condition with **udev**. As a consequence, the **mkrescue** operation failed with the following message:

```
ERROR: BUG BUG BUG! Could not determine size of disk
```

Therefore it was not possible to create the rescue image. The bug has been fixed, and rescue image creation now works as expected. (BZ#[1388653](#))

ReaR no longer requires dosfsck and efibootmgr on non-UEFI systems

Previously, **ReaR** (Relax-and-Recover) incorrectly required the **dosfsck** and **efibootmgr** utilities installed on systems that do not use UEFI. As a consequence, if the utilities were missing, the **rear mkrescue** command failed with an error. This bug has been fixed, and **ReaR** now requires the mentioned utilities to be installed only on UEFI systems. (BZ#[1479002](#))

ReaR no longer fails with NetBackup and has more reliable network configuration

Previously, two problems in the startup procedure of the rescue system caused the **ReaR** (Relax-and-Recover) restore process to fail when using the **NetBackup** method. The system's init scripts were sourced instead of executed when used by **ReaR**. As a consequence, the **NetBackup** init script aborted the system-setup process. Additionally, processes created by the system setup were immediately terminated. This affected the **dhclient** tool as well, and in some cases caused an IP address conflict. With this update, both bugs have been fixed. As a result, **ReaR** works properly with the **NetBackup** method, and network configuration using DHCP is more reliable. (BZ#[1506231](#))

ReaR recovery no longer fails when backup integrity checking is enabled

Previously, if **ReaR** (Relax-and-Recover) was configured to use backup integrity checking (**BACKUP_INTEGRITY_CHECK=1**), the recovery process always failed because the **md5sum** command could not find the backup archive. This bug has been fixed, and the described problem no longer occurs. (BZ#[1532676](#))

CHAPTER 35. STORAGE

DM Multipath no longer crashes when adding a feature to an empty string

Previously, the DM Multipath service terminated unexpectedly when it attempted to add a feature to the features string of a built-in device configuration that had no features string. With this update, DM Multipath first checks if the features string exists, and creates one if necessary. As a result, DM Multipath no longer crashes when trying to modify a nonexistent features string. (BZ#[1459370](#))

I/O operations no longer hang with RAID1

Previously, the kernel did not handle Multiple Devices (MD) I/O errors properly in **dm-raid**. As a consequence, the I/O sometimes became unresponsive. With this update, **dm-raid** now handles I/O errors correctly, and I/O operations no longer hang with RAID1. (BZ#1506338)

CHAPTER 36. SYSTEM AND SUBSCRIPTION MANAGEMENT

Yum no longer crashes in certain nss and nspr update scenario

Previously, when the **yum** installer updated a certain combination of nss and nspr package versions, the transaction sometimes terminated prematurely due to a following symbol lookup error:

```
/lib64/libnsssysinit.so: undefined symbol: PR_GetEnvSecure
```

This then caused stale rpm locks. **Yum** has been updated to correctly deal with this particular nss and nspr update scenario. As a result, **yum** does not terminate anymore in the described scenario. (BZ#[1458841](#))

The fastestmirror plug-in now orders mirrors before the metadata download

Previously, when the **yum** installer ran for the first time after a cache cleanup, the **fastestmirror** plug-in did not select the fastest mirror before metadata download. This sometimes caused a delay if some mirrors were slow or unavailable. With this update, the **fastestmirror** plug-in has been modified to have effect on mirror selection before metadata download. As a result, the mirrors are polled and arranged before metadata download, which prevents such delays. (BZ#[1428210](#))

The package-cleanup script no longer removes package dependencies of non-duplicates

Previously, running the **package-cleanup** script with the **--cleandupes** option also removed packages that depended on duplicates. Consequently, some packages were removed unintentionally. With this update, the **package-cleanup** script has been fixed to skip package dependencies of non-duplicates. Instead, the **package-cleanup** script prints a warning with a suggestion of a workaround. (BZ#[1455318](#))

rhnsd.pid is now writable only by the owner

In Red Hat Enterprise Linux 7.4, the default permissions of the **/var/run/rhnsd.pid** file were changed to **-rw-rw-rw-..**. This setting was not secure. With this update, the change has been reverted, and the default permissions of **/var/run/rhnsd.pid** are now **-rw-r--r--..** (BZ#[1480306](#))

rhncore6 now correctly reports system reboots to Satellite

Previously, if a system reboot of a Satellite client occurred during a **rhncore6** run, **rhncore6** did not report its termination to Satellite. Consequently, the status of **rhncore6** in Satellite did not update. With this update, this incorrect behavior is fixed and **rhncore6** now handles system reboots and reports the correct status to Satellite. (BZ#[1494389](#))

The rpm rhnlib -qi command now refers to the current upstream project website

Previously, the **RPM** information of the **rhnlib** package incorrectly referred to a deprecated upstream project website. With this update, the **rpm rhnlib -qi** command displays the URL of the current upstream project website. (BZ#[1503953](#))

Kernel installations using rhnsd complete successfully

If a kernel installation scheduled by the kernel was run using the **Red Hat Network Daemon** (**rhnsd**), the installation of the kernel sometimes stopped before completion. This issue has been fixed and kernel installations using **rhnsd** now complete successfully. (BZ#[1475039](#))

rhncore6 no longer modifies permissions on files in /var/cache/yum/

Previously, when the **Red Hat Network Daemon** (**rhnsd**) executed the **rhncore6** command, the command modified permissions on the files in the **/var/cache/yum/** directory incorrectly, resulting in a vulnerability. This bug has been fixed and **rhncore6** no longer modifies permissions on the files in

the `/var/cache/yum/` directory. (BZ#[1489989](#))

subscription-manager reports an RPM package if its vendor contains non-UTF8 characters

Previously, the **subscription-manager** utility assumed UTF-8 data in the **RPM** package vendor field. Consequently, if an **RPM** installed on the system contained a vendor with non-UTF8 characters, the **subscription-manager** failed to report the packages. With this update, the **subscription-manager** has been updated to ignore encoding issues in the **RPM** package vendor field. As a result, **subscription-manager** reports a package profile correctly even if the installed **RPM** has a non-UTF8 vendor. (BZ#[1519512](#))

subscription-manager now works with proxies that expect the Host header

Previously, the **subscription-manager** utility was not compatible with proxies that expect the **Host** header because it did not include the **Host** header when connecting. With this update, **subscription-manager** includes the **Host** header when connecting and is compatible with these proxies. (BZ#[1507158](#))

subscription-manager assigns valid IPv4 addresses to network.ipv4_address even if initial DNS resolution fails

Previously, when the **subscription-manager** utility failed to resolve the IPv4 address of a system, it incorrectly assigned the loopback interface address **127.0.0.1** for the **network.ipv4_address** fact. This occurred even when there was a valid interface with a valid IP address. With this update, if **subscription-manager** fails to resolve the IPv4 address of a system, it gathers IPv4 addresses from all interfaces except the loopback interface and assigns the valid IPv4 addresses for the **network.ipv4_address** fact. (BZ#[1476817](#))

virt-who ensures that provided options fit the same virtualization type

With this update, the **virt-who** utility ensures that all command-line options provided by the user are compatible with the intended virtualization type. In addition, if **virt-who** detects an incompatible option, it provides a corresponding error message. (BZ#[1461417](#))

virt-who configuration no longer resets on upgrade or reinstall

Previously, upgrading or reinstalling **virt-who** reset the configuration of the `/etc/virt-who.conf` file to default values. This update changes the packaging of **virt-who** to prevent overwriting configuration files, which ensures the described problem no longer occurs. (BZ#[1485865](#))

virt-who now reads the 'address' field provided by RHEVM to discover and report the correct host name

Previously, if the **virt-who** utility reported on a Red Hat Virtualization (RHV) host and the **hypervisor_id=hostname** option was used, **virt-who** displayed an incorrect host name value. This update ensures that **virt-who** reads the correct field value in the described circumstances and as a result, the proper host name is displayed. (BZ#[1389729](#))

CHAPTER 37. VIRTUALIZATION

Guests no longer shut down unexpectedly during reboot

On a Red Hat Enterprise Linux 7.4 guest running on **qemu-kvm-1.5.3-139.el7**, if the **i6300esb watchdog** was set to **poweroff**, the watchdog was triggered when shutting down due to the timeout being calculated incorrectly. Consequently, when rebooting the guest, it shut down instead. With this update, the timeout calculations in **qemu-kvm** have been corrected. As a result, the virtual machine reboots properly. (BZ#[1470244](#))

Guests accessed using a serial console no longer become unresponsive

Previously, if a client opened a host-side pseudoterminal device (pty) on a KVM guest pty serial console and did not read from it, the guest in some cases became unresponsive because of blocking read/write calls. With this update, the host-side pty open mode was set to non-blocking. As a result, the guest machine does not become unresponsive in the described scenario. (BZ#[1455451](#))

virt-v2v now warns about not converting PCI passthrough devices

The **virt-v2v** utility currently cannot convert PCI passthrough devices and thus ignores them in the conversion process. Prior to this update, however, attempting to convert a guest virtual machine with a PCI passthrough device successfully converted the guest, but did not provide any warning about the ignored PCI passthrough device. Now, converting such a guest logs an appropriate warning message during the conversion. (BZ#[1472719](#))

When importing OVAs, virt-v2v now parses MAC addresses

Previously, the **virt-v2v** utility did not parse the MAC addresses of network interfaces when importing Open Virtual Appliances (OVAs). Consequently, the converted guest virtual machines had network interfaces with different MAC addresses, resulting in the network setup breaking. With this release, **virt-v2v** parses the MAC addresses, if available, of network interfaces when importing OVAs. As a result, network converted guests have the same MAC addresses as specified in the OVAs and the network setup does not break. (BZ#[1506572](#))

PART III. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 7.5.

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

CHAPTER 38. 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 39. 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 40. CLUSTERING

The pcs tool now manages bundle resources in Pacemaker

As a Technology Preview starting with Red Hat Enterprise Linux 7.4, the **pcs** tool supports bundle resources. You can now use the **pcs resource bundle create** and the **pcs resource bundle update** commands to create and modify a bundle. You can add a resource to an existing bundle with the **pcs resource create** command. For information on the parameters you can set for a **bundle** resource, run the **pcs resource bundle --help** command. (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](#))

CHAPTER 41. DESKTOP

wayland available as a Technology Preview

The **wayland** display server protocol is now available in Red Hat Enterprise Linux as a Technology Preview. This update adds 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 impossible.
- 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 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 42. 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)

pNFS SCSI layout is now available for client and server

Client and server support for parallel NFS (pNFS) SCSI layouts is provided as a Technology Preview starting with Red Hat Enterprise Linux 7.3. 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)

OverlayFS

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

OverlayFS remains a Technology Preview in Red Hat Enterprise Linux 7.5 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 as of Red Hat Enterprise Linux 7.5 release. For details, see **Non-standard behavior** in the **Documentation/filesystems/overlayfs.txt** file. (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)

New package: ima-evm-utils

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.

The ima-evm-utils package is now available as a Technology Preview. (BZ#1384450)

CHAPTER 43. 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)

ibmvnic Device Driver

Starting with Red Hat Enterprise Linux 7.3, the **ibmvnic** Device Driver has been available as a Technology Preview for IBM POWER architectures. vNIC (Virtual Network Interface Controller) 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. (BZ#1391561, BZ#947163)

CHAPTER 44. 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 7.5, the **criu** packages have been upgraded to upstream version 3.5, which provides a number of bug fixes and enhancements. In addition, support for IBM z Systems and the 64-bit ARM architecture has been added. (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

As a Technology Preview, this update adds the **kexec fast reboot** feature, which 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)

Unprivileged access to name spaces can be enabled as a Technology Preview

You can now set the **namespace.unpriv_enable** kernel command-line option if required, as a Technology Preview.

The default setting is off.

When set to **1**, issuing a call to the **clone()** function with the flag **CLONE_NEWNS** as an unprivileged user no longer returns an error and allows the operation.

However, to enable the unprivileged access to name spaces, the **CAP_SYS_ADMIN** flag has to be set in some user name space to create a mount name space. (BZ#1350553)

SCSI-MQ as a Technology Preview in the q1a2xxx driver

The **qla2xxx** driver updated in Red Hat Enterprise Linux 7.4 can now enable the use of SCSI-MQ (multiqueue) with the **ql2xmqsupport=1** module parameter. The default value is **0** (disabled). The SCSI-MQ functionality is provided as a Technology Preview when used with the **qla2xxx** driver.

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. A fix is being tested but was not ready in time for Red Hat Enterprise Linux 7.4 General Availability. (BZ#1414957)

NVMe over Fibre Channel is now available as a Technology Preview

The NVMe over Fibre Channel transport type is now available as a Technology Preview. 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 over Fibre Channel in the **lpfc** driver, edit the **/etc/modprobe.d/lpfc.conf** file and add one or both of the following options:

- To enable the NVMe mode of operation, add the **lpfc_enable_fc4_type=3** option.
- To enable target mode, add the **lpfc_enable_nvmet=<wwpn list>** option, where **<wwpn list>** is a comma-separated list of World-Wide Port Name (WWPN) values with the **0x** prefix.

To configure an NVMe target, use the **nvmectl** utility.

NVMe over Fibre Channel 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.

In Red Hat Enterprise Linux 7.5, NVMe over Fibre Channel is available only with Broadcom 32Gbit adapters, which use the **lpfc** driver. (BZ#[1387768](#), BZ#1454386)

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

With this update, **perf cqm** has been replaced by the approach based on the **resctrl** file system, which address all of the aforementioned problems. (BZ#[1457533](#), BZ#1288964)

CHAPTER 45. REAL-TIME KERNEL

The **SCHED_DEADLINE** scheduler class as Technology Preview

The **SCHED_DEADLINE** scheduler class for the real-time kernel, which was introduced in Red Hat Enterprise Linux 7.4, continues to be available as a Technology Preview. The scheduler enables predictable task scheduling based on application deadlines. **SCHED_DEADLINE** benefits periodic workloads by reducing application timer manipulation. (BZ#1297061)

CHAPTER 46. 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 qlcnict driver

Support for Single-Root I/O virtualization (SR-IOV) has been added to the `qlcnict` 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 `qlcnict` driver remains fully supported. (BZ#1259547)

The libnftnl and nftables packages

Starting with Red Hat Enterprise Linux 7.3., the `nftables` and `libnftnl` packages are available as a Technology Preview.

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's API over the **`libmnl`** library. (BZ#1332585)

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 47. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE

Red Hat Enterprise Linux System Roles

Red Hat Enterprise Linux System Roles, available as a Technology Preview, is 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>. (BZ#1313263)

CHAPTER 48. 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 49. 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)

CHAPTER 50. VIRTUALIZATION

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)

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 7.5.

CHAPTER 51. NEW DRIVERS

Storage Drivers

- USB Type-C Connector Class (typec.ko.xz):
- USB Type-C Connector System Software Interface driver (typec_ucsi.ko.xz):
- TCM QLA2XXX series NPIV enabled fabric driver (tcm_qla2xxx.ko.xz):
- Chelsio FCoE driver (csiostor.ko.xz): 1.0.0-ko

Network Drivers

- Software simulator of 802.11 radio(s) for mac80211 (mac80211_hwsim.ko.xz):
- Vsock monitoring device. Based on nlmon device. (vsockmon.ko.xz):
- Cavium LiquidIO Intelligent Server Adapter Virtual Function Driver (liquidio_vf.ko.xz): 1.6.1
- Cavium LiquidIO Intelligent Server Adapter Driver (liquidio.ko.xz): 1.6.1
- Mellanox firmware flash lib (mlxfw.ko.xz):
- Intel OPA Virtual Network driver (opa_vnic.ko.xz):
- Broadcom NetXtreme-C/E RoCE Driver Driver (bnxt_re.ko.xz):
- VMware Paravirtual RDMA driver (vmw_pvrDMA.ko.xz):

Graphics Drivers and Miscellaneous Drivers

- MC Driver for Intel SoC using Pondicherry memory controller (pnd2_edac.ko.xz):
- ALPS HID driver (hid-alps.ko.xz):
- Intel Corporation DAX device (device_dax.ko.xz):
- Synopsys DesignWare DMA Controller platform driver (dw_dmac.ko.xz):
- Synopsys DesignWare DMA Controller core driver (dw_dmac_core.ko.xz):
- Intel SunrisePoint PCH pinctrl/GPIO driver (pinctrl-sunrisepoint.ko.xz):
- Intel Lewisburg pinctrl/GPIO driver (pinctrl-lewisburg.ko.xz):
- Intel Cannon Lake PCH pinctrl/GPIO driver (pinctrl-cannonlake.ko.xz):
- Intel Denverton SoC pinctrl/GPIO driver (pinctrl-denverton.ko.xz):
- Intel Gemini Lake SoC pinctrl/GPIO driver (pinctrl-geminilake.ko.xz):
- Intel pinctrl/GPIO core driver (pinctrl-intel.ko.xz):

CHAPTER 52. UPDATED DRIVERS

Storage Driver Updates

- The QLogic Fibre Channel HBA driver (qla2xxx.ko.xz) has been updated to version 9.00.00.00.07.5-k1.
- The Cisco FCoE HBA Driver driver (fnic.ko.xz) has been updated to version 1.6.0.34.
- The Emulex OneConnectOpen-iSCSI driver (be2iscsi.ko.xz) has been updated to version 11.4.0.1.
- The QLogic FCoE driver (bnx2fc.ko.xz) has been updated to version 2.11.8.
- The Microsemi Smart Family Controller driver (smartpqi.ko.xz) has been updated to version 1.1.2-126.
- The Emulex LightPulse Fibre Channel SCSI driver (lpfc.ko.xz) has been updated to version 0:11.4.0.4.
- The LSI MPT Fusion SAS 3.0 Device driver (mpt3sas.ko.xz) has been updated to version 16.100.00.00.
- The QLogic QEDF 25/40/50/100Gb FCoE driver (qedf.ko.xz) has been updated to version 8.20.5.0.
- The Avago MegaRAID SAS driver (megaraid_sas.ko.xz) has been updated to version 07.702.06.00-rh2.
- The HP Smart Array Controller driver (hpsa.ko.xz) has been updated to version 3.4.20-0-RH2.

Network Driver Updates

- The Realtek RTL8152/RTL8153 Based USB Ethernet Adapters driver (r8152.ko.xz) has been updated to version v1.08.9.
- The Intel(R) 10 Gigabit PCI Express Network driver (ixgbe.ko.xz) has been updated to version 5.1.0-k-rh7.5.
- The Intel(R) Ethernet Switch Host Interface driver (fm10k.ko.xz) has been updated to version 0.21.7-k.
- The Intel(R) Ethernet Connection XL710 Network driver (i40e.ko.xz) has been updated to version 2.1.14-k.
- The Intel(R) 10 Gigabit Virtual Function Network driver (ixgbev.ko.xz) has been updated to version 4.1.0-k-rh7.5.
- The Intel(R) XL710 X710 Virtual Function Network driver (i40evf.ko.xz) has been updated to version 3.0.1-k.
- The Elastic Network Adapter (ENA) driver (ena.ko.xz) has been updated to version 1.2.0k.
- The Cisco VIC Ethernet NIC driver (enic.ko.xz) has been updated to version 2.3.0.42.
- The Broadcom BCM573xx network driver (bnxt_en.ko.xz) has been updated to version 1.8.0.

- The QLogic FastLinQ 4xxxx Core Module driver (qed.ko.xz) has been updated to version 8.10.11.21.
- The QLogic 1/10 GbE Converged/Intelligent Ethernet driver (qlcnict.ko.xz) has been updated to version 5.3.66.
- The Mellanox ConnectX HCA Ethernet driver (mlx4_en.ko.xz) has been updated to version 4.0-0.
- The Mellanox ConnectX HCA low-level driver (mlx4_core.ko.xz) has been updated to version 4.0-0.
- The Mellanox Connect-IB, ConnectX-4 core driver (mlx5_core.ko.xz) has been updated to version 5.0-0.

Graphics Driver and Miscellaneous Driver Updates

- The standalone VMware SVGA device drm driver (vmwgfx.ko.xz) has been updated to version 2.14.0.0.

CHAPTER 53. 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 7.5.

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`

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>
<code>ypserv</code> <code>ypbind</code> <code>portmap</code> <code>yp-tools</code>	Identity Management in Red Hat Enterprise Linux

Use the Go Toolset instead of `golang`

The `golang` package has been updated to version 1.9 with Red Hat Enterprise Linux 7.5.

The `golang` package, available in the Optional channel, will be removed from a future minor release of Red Hat Enterprise Linux 7. Developers are encouraged to use the **Go Toolset** instead, which is currently available as a Technology Preview 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, available as a Technology Preview, 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
- 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

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 Ethernet adapter controlled by the **be2net** driver has been deprecated:
 - TIGERSHARK NIC, PCI ID 0x0700
- The following adapters from the **be2iscsi** driver have been deprecated:
 - Emulex OneConnect 10Gb iSCSI Initiator (generic), PCI ID 0x212
 - OCe10101, OCm10101, OCe10102, OCm10102 BE2 adapter family, PCI ID 0x702
 - OCe10100 BE2 adapter family, PCI ID 0x703
- The following adapters from the **lpfc** driver have been deprecated:
 - BladeEngine 2 (BE2) Devices
 - TIGERSHARK FCOE, PCI ID 0x0704
 - Fibre Channel (FC) Devices
 - FIREFLY, PCI ID 0x1ae5
 - PROTEUS_VF, PCI ID 0xe100
 - BALIUS, PCI ID 0xe131
 - PROTEUS_PF, PCI ID 0xe180
 - RFLY, PCI ID 0xf095
 - PFLY, PCI ID 0xf098
 - LP101, PCI ID 0xf0a1
 - TFLY, PCI ID 0xf0a5
 - BSMB, PCI ID 0xf0d1
 - BMID, PCI ID 0xf0d5
 - ZSMB, PCI ID 0xf0e1
 - ZMID, PCI ID 0xf0e5
 - NEPTUNE, PCI ID 0xf0f5
 - NEPTUNE_SCSP, PCI ID 0xf0f6
 - NEPTUNE_DCSP, PCI ID 0xf0f7
 - FALCON, PCI ID 0xf180
 - SUPERFLY, PCI ID 0xf700
 - DRAGONFLY, PCI ID 0xf800

- CENTAUR, PCI ID 0xf900
- PEGASUS, PCI ID 0xf980
- THOR, PCI ID 0xfa00
- VIPER, PCI ID 0xfb00
- LP10000S, PCI ID 0xfc00
- LP11000S, PCI ID 0xfc10
- LPE11000S, PCI ID 0xfc20
- PROTEUS_S, PCI ID 0xfc50
- HELIOS, PCI ID 0xfd00
- HELIOS_SCSP, PCI ID 0xfd11
- HELIOS_DCSP, PCI ID 0xfd12
- ZEPHYR, PCI ID 0xfe00
- HORNET, PCI ID 0xfe05
- ZEPHYR_SCSP, PCI ID 0xfe11
- ZEPHYR_DCSP, PCI ID 0xfe12

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](#).

PART V. KNOWN ISSUES

This part documents known problems in Red Hat Enterprise Linux 7.5.

CHAPTER 54. AUTHENTICATION AND INTEROPERABILITY

A crash is reported after an unsuccessful lightweight CA key retrieval

When using Identity Management (IdM), if retrieving the lightweight certificate authority (CA) key fails for some reason, the operation terminates unexpectedly with an uncaught exception. The exception results in a crash report. (BZ#[1478366](#))

OpenLDAP causes programs to fail immediately in case of incorrect configuration

Previously, the Mozilla implementation of Network Security Services (Mozilla NSS) silently ignored certain misconfigurations in the OpenLDAP suite, which caused programs to fail only on connection establishment. With this update, OpenLDAP has switched from Mozilla NSS to OpenSSL (see the release note for BZ#[1400578](#)). With OpenSSL, the TLS context is established immediately, and therefore programs fail immediately. This behavior prevents potential security risks, such as keeping non-working TLS ports open.

To work around this problem, verify and fix your OpenLDAP configuration. (BZ#[1515833](#))

OpenLDAP reports failures when CACertFile or CACertDir point to an invalid location

Previously, if the CACertFile or CACertDir options pointed to an unreadable or otherwise unloadable location, the Mozilla implementation of Network Security Services (Mozilla NSS) did not necessarily consider it a misconfiguration. With this update, the OpenLDAP suite has switched from Mozilla NSS to OpenSSL (see the release note for BZ#[1400578](#)). With OpenSSL, if CACertFile or CACertDir point to such an invalid location, the problem is no longer silently ignored.

To avoid the failures, remove the misconfigured option, or make sure it points to a loadable location.

Additionally, OpenLDAP now applies stricter rules for the contents of the directory to which CACertDir points. If you experience errors when using certificates in this directory, it is possible the directory is in an inconsistent state. To fix this problem, run the **openssl rehash** command on the folder.

For details on CACertFile and CACertDir, see these man pages: `ldap.conf(5)`, `slapd.conf(5)`, `slapd-config(5)`, and `ldap_set_option(3)`. (BZ#[1515918](#), BZ#[1515839](#))

OpenLDAP does not update TLS configuration after inconsistent changes in `cn=config`

With this update, OpenLDAP has switched from the Mozilla implementation of Network Security Services (Mozilla NSS) to OpenSSL (see the release note for BZ#[1400578](#)). With OpenSSL, inconsistent changes of the TLS configuration in the `cn=config` database break the TLS protocol on the server, and configuration is not updated as expected. To avoid this problem, use only one change record to update the TLS configuration in `cn=config`. See the `ldif(5)` man page for a definition of a change record. (BZ#[1524193](#))

Identity Management terminates connections unexpectedly

Due to a bug in Directory Server, Identity Management (IdM) terminates connections unexpectedly after a certain amount of time, and authentication fails with the following error:

```
kinit: Generic error (see e-text) while getting initial credentials
```

The problem occurs if you installed IdM on Red Hat Enterprise Linux 7.5 from an offline media. To work around the problem, run **yum update** to receive the updated 389-ds-base package which fixes the problem. (BZ#[1544477](#))

Directory Server can terminate unexpectedly during shutdown

Directory Server uses the **nunc-stans** framework to manage connection events. If a connection is closed when shutting down the server, a **nunc-stans** job can access a freed connection structure. As a consequence, Directory Server can terminate unexpectedly. Because this situation occurs in a late state of the shutdown process, data is not corrupted or lost. Currently, no workaround is available.
(BZ#[1517383](#))

CHAPTER 55. CLUSTERING

Data corruption occurs on RAID 10 reshape on top of VDO with el7 kernel.

RAID 10 reshape (with both LVM and `mdadm`) on top of VDO corrupts data and can eventually trigger the `raid10.c:1011` kernel bug. Stacking RAID 10 (or other RAID types) on top of VDO does not take advantage of the deduplication/compression capabilities of VDO and is not recommended. (BZ#[1528466](#), BZ#1530776)

CHAPTER 56. COMPILER AND TOOLS

Memory consumption of applications using libcurl grows with each TLS connection

The **Network Security Services** (NSS) `PK11_DestroyGenericObject()` function does not release resources allocated by `PK11_CreateGenericObject()` early enough. Consequently, the memory allocated by applications using the libcurl package can grow with each TLS connection.

To work around this problem:

- Re-use existing TLS connections where possible or
- Use certificates and keys from the **NSS** database instead of loading them from files directly using libcurl (BZ#[1510247](#))

oProfile and perf can not sample events on 2nd generation Intel Xeon Phi processors when NMI watchdog is disabled

Due to a performance counter hardware error, sampling performance events with the default hardware event `CPU_CLK_UNHALTED` may fail on 2nd generation Intel Xeon Phi processors. As a consequence, the **oProfile** and **perf** tools fail to receive any samples when the NMI watchdog is disabled. To work around this problem, enable NMI watchdog before running the **perf** or **operf** command:

```
echo 1 > /proc/sys/kernel/nmi_watchdog
...
operf some_examined_program
opreport
...
```

Note that this workaround allows only the selected tool to work correctly, but not the NMI watchdog, because it is based on the NMI watchdog using the erroneous counter. (BZ#1536004)

CHAPTER 57. DESKTOP

Cannot install downloaded RPM files from **Nautilus**

The **yum** backend to **PackageKit** does not support getting details about local files. As a consequence, when an RPM file is double clicked in the **Nautilus** file manger, the file is not installed, and the following error message is returned:

```
Sorry, this did not work, File is not supported
```

To work around this problem, either install the **gnome-packagekit** package to handle the double-click action, or manually install the files using the **yum** utility. (BZ#[1434477](#))

Caps Lock LED status

When using an UTF-8 keymap, even though the caps lock function works properly, the caps lock LED is not updated while in TTY mode. For the LED to be correctly updated, starting from Red Hat Enterprise Linux 7.5, the administrator needs to create the `/etc/udev/rules.d/99-kbd.rules` configuration file as follows:

```
ACTION=="add", SUBSYSTEM=="leds",  
ENV{DEVPATH}=="*/input*::capslock",  
ATTR{trigger}="kbd-ctrl1lock"
```

To reload the new udev rule, run these commands:

```
# udevadm control --reload-rules  
# udevadm trigger
```

After this change, when pressing the caps lock key, caps lock LED changes its status as expected. (BZ#[1470932](#), BZ#1256895)

Inconsistent **GNOME Shell** versions

The GNOME desktop environment currently displays different versions of **GNOME Shell**. For example, the version returned by the **gnome-shell --version** command is different from the version found in the **Details** section of **Settings**. (BZ#[1511454](#))

Uninstall the 32-bit version of **flatpak**

Users are advised to uninstall the 32-bit version of the flatpak packages before updating to Red Hat Enterprise Linux 7.5 to prevent possible multilib conflicts. (BZ#1512940)

GNOME downgrade does not work

With the new version of GNOME (3.22) introduced in Red Hat Enterprise Linux 7.4, downgrading GNOME from version 3.22 to 3.14 using the **yum downgrade** or **dnf downgrade** commands is no longer possible. The only workaround lies in replacing the GNOME-related packages with their old versions. If you decide to downgrade manually, read the GNOME 3.16-3.22 release notes to find which functionalities you are losing. (BZ#[1451876](#))

wayland ignores keyboard grabs issued by **X11** applications, such as virtual machines viewers

Currently, when running through the **Xwayland** server, graphical clients that rely on the **X11** software, such as remote desktop viewers or virtual machine managers, are unable to obtain the system keyboard shortcuts for their own use. As a consequence, activating these shortcuts in a guest window, such as a **virt-manager** guest display, affects the local desktop instead of the guest.

To work around the problem, use a **Wayland** native client with support for Wayland shortcuts inhibitor protocol, or switch back to the default GNOME session on **X11** to run the **X11** clients that require system keyboard shortcuts.

Note that Wayland is available as a Technology Preview. (BZ#1500397)

Superuser should not run graphical sessions

Opening a graphical session for the root user causes various bugs. The reason is that a graphical session is not meant to be used by superuser as it can cause serious and unexpected issues, is non-secure, and is against Unix principles. (BZ#1539772)

Keyboard not working in VM browsed by **remote-viewer** and **virt-viewer**

When run inside a Wayland session, **remote-viewer** and **virt-viewer** utilities do not recognize key events in a virtual machine. Moreover, Xwayland reports the following error:

```
send_key: assertion 'scancode != 0'
```

(BZ#1540056)

gnome-system-log does not work on wayland

Currently, when logged in a **Wayland** session, the root user is not allowed to access the user's Xwayland display. As a consequence, running the **gnome-system-log** utility in terminal does not display system log files.

To work around this problem, run the following **xhost** server access control program as follows:

```
$ xhost +si:localuser:root
```

(BZ#1537529)

GUI screen is shown incorrectly

The X driver for Emulex Pilot2 and Pilot3 cards contains a bug when running at depth of color 16. This bug makes the graphics display unusable at this depth.

To make the display usable in some configurations, use 24 bpp image format. Alternatively, disable the shadow framebuffer abstraction layer in the **xorg.conf** file by using the **ShadowFB off** option. Note that disabling the shadow framebuffer may have significant performance impact. (BZ#1499129)

xrandr fails to provide some video modes

Different video drivers for **X11** have different heuristics for adding display resolutions. In particular, the Intel and generic modesetting drivers provide different sets of video modes for some laptop displays. Consequently, some non-native video modes may not be available in all configurations.

To work around this problem, use a different video driver, or add resolutions to the output manually using the **xrandr(1)** command-line utility. (BZ#1478625)

radeon fails to reset hardware correctly

The **radeon** kernel driver currently does not reset hardware in the **kexec** context correctly. Instead, **radeon** falls over, which causes the rest of the **kdump** service to fail.

To work around this bug, blacklist **radeon** in **kdump** by adding the following line to the **/etc/kdump.conf** file:

```
dracut_args --omit-drivers "radeon"
```

```
force_rebuild 1
```

Restart the machine and **kdump**. After starting **kdump**, the **force_rebuild 1** line may be removed from the configuration file.

Note that in this scenario, no graphics will be available during **kdump**, but **kdump** will complete successfully. (BZ#1509444)

nouveau fails to load Nvidia secboot firmware

In some Dell Coffeelake systems, the **nouveau** kernel module fails to load Nvidia secboot firmware for the pascal cards. As a consequence, Nvidia GPU on these systems occasionally does not work, and some of the Display ports on the system thus do not work as well.

If this bug causes trouble booting, blacklist **nouveau** to mitigate the problem. Note that this, however, will not make non-functional ports on the machine work correctly. (BZ#1535168)

xchat status icon disappears from Top Icons panel

The **Xchat** status icon indicating incoming personal messages disappears from top icons panel after suspending the system and resuming it again.

Top icons installed using **Gnome Software** preserve the suspend mode and do not disappear from the panel. (BZ#1544840)

GDM does not activate hotplugged monitors

When a machine is booted without a monitor connected, the **GNOME Display Manager (GDM)** screen remains deactivated when a monitor is plugged in.

As a workaround, kill **GDM** while the monitor is plugged in by running:

```
# systemctl restart gdm.service
```

Alternatively, use the **xrandr** utility to activate the monitor. (BZ#1497303)

Wacom Expresskeys Remote not detected as tablet

The **gnome-shell** and **control-center** utilities do not detect unpaired **Wacom Expresskeys Remote** devices (EKR). As a consequence, within the Wacom settings, there is no way to map the buttons on the **EKR**.

Currently, **EKR** works only when it is paired to a tablet with a built-in pad. (BZ#1543631)

Synaptics dependency removes xorg-x11-drivers

Later releases of Red Hat Enterprise Linux 7 contain the **xorg-x11-drv-libinput** driver for X, which can potentially provide a superior experience for some input devices. Users attempting to switch to **xorg-x11-drv-libinput** can try removing the **xorg-x11-drv-synaptics** driver, which is required by the **xorg-x11-drivers** package. However, removing **synaptics** requires removing **xorg-x11-drivers**.

To work around this issue, remove **xorg-x11-drivers**. This package exists only to install a reasonable collection of drivers at system setup time, and removing it has no runtime impact. Any X driver already installed will be updated as expected. (BZ#1516970)

T470s docking station jack does not work on resume

After suspending and resuming ThinkPad T470s connected to the docking station with analog audio input or output, the user does not receive any output sound. This problem does not affect the analog audio input or output in the ThinkPad laptop. (BZ#1548055)

Screen occasionally turns off when `xrandr` is executed

With the **Nouveau** driver, RANDR operations combined with heavy 3D load, such as querying the screen resolution, may cause screen flickering.

Flickering can be avoided by minimizing concurrent 3D and RANDR operations. Hence, query or resize the screen while 3D usage is minimal. (BZ#[1545550](#))

HDMI and DP for 8th generation Intel Core processors not enumerating sound inputs

In Red Hat Enterprise Linux, support for **alpha** status hardware is disabled in the i915 driver by default. which causes that i915 never binds to the audio driver. As a consequence, HDMI and DP video and audio standards for 8th generation Intel Core processors do not enumerate sound inputs.

To work around this issue, boot your system with the **i915.alpha_support=1** line added to the kernel command line. (BZ#[1540643](#))

Tray icons are non-responsive for auto-started applications

The **GNOME Shell TopIcons** extension, which shows legacy tray icons on the top of the screen, does not work for auto-started applications: the tray icons are non-responsive. This bug does not include applications started after the GNOME Session starts.

As a workaround, follow this short procedure to restart the GNOME session: 1. press **Alt + F2**, 2. type **r**, 3. press **Enter**. (BZ#[1550115](#))

Inconsistent panel color on login screen

When logging to a **GNOME Classic** session, suspending the laptop and resuming it again, the top panel on login screen is white, instead of black.

This problem does not affect **GNOME Classic** functionality. (BZ#[1541021](#))

Additional displays are mirrored after attaching a VM guest

When opening a guest VM monitor and enabling an additional display from the **remote-viewer** menu, the content of the first display is mirrored to the newly attached one.

As a workaround, resize the **remote-viewer** frame of any display. The desktop environment will be extended to both displays and guest displays will be properly rearranged. (BZ#[1539686](#))

CHAPTER 58. INSTALLATION AND BOOTING

Selecting the Lithuanian language causes the installer to crash

If you select the Lithuanian (Lietuvių) language on the first screen of the graphical installer and press **Continue** (Tęsti), the installer crashes and displays a traceback message. To work around this problem, either use a different language, or avoid the graphical installer and use a different approach such as the text mode or a Kickstart installation. (BZ#[1527319](#))

oscap-anaconda-addon fails to remediate when installing in TUI using Kickstart

The **OpenSCAP Anaconda** add-on fails to fully remediate a machine to the specified security policy when the system is installed using a Kickstart file that sets installation display mode to the text-based user interface (TUI) using the **text** Kickstart command. The problem occurs because packages required for the remediation are not installed.

To work around this problem, you can either use the graphical installer or add packages required by the security policy to the **%packages** section of the Kickstart file manually. (BZ#[1547609](#))

The grub2-mkimage command fails on UEFI systems by default

The **grub2-mkimage** command may fail on UEFI systems with the following error message:

```
error: cannot open `usr/lib/grub/x86_64-efi/moddep.lst': No such file or directory.
```

This error is caused by the package **grub2-efi-x64-modules** package missing from the system. The package is missing due to a known issue where it is not part of the default installation, and it is not marked as a dependency for **grub2-tools** which provides the **grub2-mkimage** command.

The error also causes some other tools which depend on it, such as **ReaR**, to fail.

To work around this problem, install the **grub2-efi-x64-modules**, either manually using **Yum**, or by adding it to the Kickstart file used for installing the system. (BZ#[1512493](#))

Kernel panic during RHEL 7.5 installation on HPE BL920s Gen9 systems

A known issue related to the fix for the Meltdown vulnerability causes a kernel panic with a NULL pointer dereference during the installation of Red Hat Enterprise Linux 7.5 on HPE BL920s Gen2 (Superdome 2) systems. When the problem appears, the following error message is displayed:

```
WARNING: CPU: 576 PID: 3924 at
kernel/workqueue.c:1518__queue_delayed_work+0x184/0x1a0
```

Then the system reboots, or enters an otherwise faulty state.

There are multiple possible workarounds for this problem:

- Add the **nopti** option to the kernel command line using the boot loader. Once the system finishes booting, upgrade to the latest RHEL 7.5 kernel.
- Install RHEL 7.4, and then upgrade to the latest RHEL 7.5 kernel.
- Install RHEL 7.5 on a single blade. Once the system is installed, upgrade to the latest RHEL 7.5 kernel, and then add additional blades as required. (BZ#[1540061](#))

The READONLY=yes option is not sufficient to configure a read-only system

In Red Hat Enterprise Linux 6, the **READONLY=yes** option in the **/etc/sysconfig/readonly-root** file was used to configure a read-only system partition. In Red Hat Enterprise Linux 7, the option is no longer sufficient, because **systemd** uses a new approach to mounting the system partition.

To configure a read-only system in Red Hat Enterprise Linux 7:

- Set the **READONLY=yes** option in **/etc/sysconfig/readonly-root**.
- Add the **ro** option to the root mount point in the **/etc/fstab** file. (BZ#[1444018](#))

CHAPTER 59. KERNEL

Security patches addressing Spectre and Meltdown issues can cause performance loss

Security patches to address issues reported in CVE-2017-5754, CVE-2017-5715, and CVE-2017-5753 have been implemented. For more information on the issues, including their impact, detection and resolution, see the Red Hat Knowledgebase article at <https://access.redhat.com/security/vulnerabilities/speculativeexecution>. The patches are enabled by default but they can cause a performance degradation.

Users can control the impact by using Red Hat Enterprise Linux Tunables. The three debugfs tunables can be enabled or disabled on the kernel command line at boot, or at runtime using debugfs controls. The tunables control Page Table Isolation (pti), Indirect Branch Restricted Speculation (ibrs), and Indirect Branch Prediction Barriers (ibpb). Red Hat enables each of the features by default as needed to protect the architecture detected at boot.

Customers who feel confident that their systems are well protected by other means and wish to disable the CVE mitigations to avoid such a performance loss, should use one of the following options:

1. Add the following flags to the kernel command line, and then reboot the kernel for the changes to take effect:

```
spectre_v2=off nopti
```

2. Run the following commands to disable the patches at runtime. The change is immediately active and does not require a reboot.

```
# echo 0 > /sys/kernel/debug/x86/pti_enabled
# echo 0 > /sys/kernel/debug/x86/retp_enabled
# echo 0 > /sys/kernel/debug/x86/ibrs_enabled
```

For more information on controlling the performance impact of the CVE mitigations, refer to the Red Hat Knowledgebase article available at <https://access.redhat.com/articles/3311301>.

See also the Diagnose tab at <https://access.redhat.com/security/vulnerabilities/speculativeexecution>. (BZ#1532547)

The KSC does not support the xz compression

The Kernel module Source Checker (the ksc tool) is unable to process the **xz** compression method, reporting the error:

```
File format not recognized (Only kernel object files are supported)
```

To work around the problem, manually uncompress any third party modules using the **xz** compression before running the **ksc** tool. (BZ#1441455)

The update of megaraid_sas can lead to a performance decrease

The **megaraid_sas** driver has been updated to version 06.811.02.00-rh1, which brings a number of performance improvements over the previous version. However, in some cases, with configurations based on Solid-state Drives (SSD) a performance decrease has been observed. To work around this problem, set the corresponding **queue_depth** parameter in the **/sys/** directory to a higher value up to 256, which brings the performance back to its original level. (BZ#1367444)

qedi fails to bind to the iSCSI PCIe function if qede is loaded

The **qede** driver, which is the ethernet driver for the QL41xxx network adapters, allocates more MSI-X vectors than needed. Consequently, the **qed** driver fails to bind to the iSCSI PCIe function exposed by the hardware. To work around this problem, unload both the **qede** and **qed** drivers, and then load only **qed**. As a result, **qed** is able to probe the iSCSI function exposed through the hardware and find any attached iSCSI targets. (BZ#1484047)

radeon causes a kernel panic

On some systems equipped with the **radeon** kernel driver as the secondary or primary GPU, the system occasionally fails to start due to a bug in the **amdgpu** graphics driver.

As a workaround, blacklist the **radeon** kernel driver. (BZ#1486100)

kdump kernel fails to boot after a CPU hot add or hot remove operation

When running Red Hat Enterprise Linux 7 on the little-endian variant of IBM Power Systems with **Kdump** enabled, the **Kdump** crashkernel will fail to boot if triggered by **kexec** after a CPU hot add or hot remove operation. To work around this problem, restart the **kdump** service after hot adding or hot removing a CPU:

```
# systemctl restart kdump.service
```

(BZ#1549355)

CHAPTER 60. 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 61. SECURITY

nss accept malformed RSA PKCS#1 v1.5 signatures made with an RSA-PSS key

The **Network Security Services** (NSS) libraries do not check the type of an RSA public key used by a server when validating signatures made using a corresponding private key. Consequently, **NSS** accept malformed RSA PKCS#1 v1.5 signatures if they are made with an RSA-PSS key. (BZ#[1510156](#))

Authentication using ssh-agent not from openssh fails

OpenSSH since version 7.4 negotiates the SHA-2 signature extension by default. Consequently, if a signature is provided by the **ssh-agent** program that is not from the current **OpenSSH** suite and that does not know the SHA-2 extension, authentication fails. To work around this problem, use the **OpenSSH ssh-agent** to provide signatures. (BZ#[1497680](#))

Parsing of openssh public keys is more strict

Previously, the parsing of public keys was changed to be more strict. As a consequence, additional spaces between the key type string and the key blob string are no longer ignored, and login attempts with such keys now fail. To work around this problem, ensure that there is only one space character between the key type and the key blob. (BZ#[1493406](#))

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 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](#))

clevis can log spurious Device is not initialized error messages

If the **Clevis** pluggable framework is in the **initramfs** image and if you have an encrypted volume configured to unlock during boot time and coincidentally you have not configured the **Clevis** binding, then the boot log shows spurious **Device is not initialized** error messages. To work around this problem, perform the **Clevis** binding step, and the error messages for the volume disappear. (BZ#[1538759](#))

Libreswan is not working properly with seccomp=enabled on all configurations

The set of allowed syscalls in the **Libreswan** SECCOMP support implementation is currently not complete. Consequently, when SECCOMP is enabled in the **ipsec.conf** file, the syscall filtering rejects even syscalls needed for proper functioning of the **pluto** daemon; the daemon is killed, and the **ipsec** service is restarted.

To work around this problem, set the **seccomp=** option back to the **disabled** state. SECCOMP support must remain disabled to run **ipsec** properly. (BZ#[1544463](#))

openSCAP RPM verification rules do not work correctly with VM and container file systems

The **rpminfo**, **rpmverify**, and **rpmverifyfile** probes do not fully support offline mode. Consequently, **OpenSCAP** RPM verification rules do not work correctly when scanning virtual machine (VM) and container file systems in offline mode.

To work around this problem, disable the RPM verification rules or perform a manual check using a guidance in the **SCAP Security Guide**. Results of scanning VM and container file systems in offline mode might contain false negatives. (BZ#[1556988](#))

Firefox and other applications using NSS become unresponsive when a smart card is inserted

The **Network Security Services** (NSS) libraries incorrectly handle smart card insertion events and states of such events. Consequently, the **Firefox** browser and other applications using **NSS** in the Gnome Display Manager (GDM) do not reliably detect the card insertion state and become unresponsive while requesting to wait for slot events.

To work around this problem, do not update the nss packages to version 3.34 and wait for the upstream version 3.36. The smart cards work correctly with the previous **NSS** version. (BZ#[1557015](#))

CHAPTER 62. SERVERS AND SERVICES

No clear indication of profile activation error in the Tuned service

Errors in the Tuned service configuration or errors occurring when loading Tuned profiles are in some cases not shown in the output of the `systemctl status tuned` command. As a consequence, if errors occur that prevent Tuned from loading, Tuned sometimes enters a state with no profile activated. To view possible error messages, consult the output of the `tuned-adm active` command and check the contents of the `/var/log/tuned/tuned.log` file. (BZ#1385838)

db_hotbackup -c should be used with caution

The `db_hotbackup` command with the `-c` option must be run by the user that owns the database. If the user is different and the log file reaches its maximal size, a new log file is created with an ownership of the user that ran the command, which consequently makes the database unusable for its owner. This note has been added to the `db_hotbackup(1)` manual page. (BZ#1460077)

Setting ListenStream= options in rpcbind.socket causes systemd-logind to fail and ssh connections to be delayed

Setting the `ListenStream=` options in the `rpcbind.socket` unit file currently causes a failure of the `systemd-logind` service and a delay in `SSH` connections that import system users from a `NIS` database. To work around the problem, remove lines with the `ListenStream=` option from `rpcbind.socket`. (BZ#1425758)

ReaR recovery process fails on non-UEFI systems with the grub2-efi-x64 package installed

Installing the `grub2-efi-x64` package, which contains the GRUB2 boot loader for UEFI systems, changes the file `/boot/grub2/grubenv` into a dead absolute symlink on systems which do not use UEFI firmware. When attempting to recover such a system using the **ReaR** (Relax and Recover) recovery tool, the process fails and the system is rendered unbootable. To work around this problem, do not install the `grub2-efi-x64` package on systems where it is not required (systems without UEFI firmware). (BZ#1498748)

ISO images generated by ReaR with Linux TSM fail to work

The password store has changed in the Linux TSM (Tivoli Storage Manager) client versions 8.1.2 and above. This means ISO images generated by **ReaR** using TSM will not work, as the TSM node password and encryption key will not be included in the ISO file. To fix this problem, add the following line into the `/etc/rear/local.conf` or `/etc/rear/site.conf` configuration file:

```
COPY_AS_IS_TSM=( /etc/adsm /opt/tivoli/tsm/client /usr/local/ibm/gsk8* )
```

(BZ#1534646)

Unexpected problems with the dbus rebase

The `dbus` package rebase with its configuration changes can cause unexpected problems. Thus, it is recommended to avoid the following actions:

- updating only the `dbus` service
- updating only parts of the system
- updating from a graphical session

On the contrary, it is recommended to reboot after executing the `yum update` command as updating several major components including `dbus` without reboot rarely works as expected. (BZ#1550582)

CHAPTER 63. STORAGE

The **kexec -e** command might cause storage errors with advanced storage controllers

When using the **kexec** utility with the **-e** option, the system does not go through the standard Linux shutdown sequence before booting the next kernel. This might cause problems for systems employing advanced storage controllers, such as the Qlogic QMH2672 Fibre Channel adapter, because these controllers rely on the shutdown sequence to assure the storage has settled at the time of a reboot. When invoking the **kexec -e** command on such systems, storage related errors might occur as the **kexec** operation progresses, and the newly loaded kernel might fail to discover some or all attached storage.

If you see similar symptoms on your system when attempting **kexec -e**, use **kexec** without the **-e** option instead. This has been observed to work reliably. (BZ#1303244)

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)

CHAPTER 64. VIRTUALIZATION

Guests reporting `cmt`, `mbmt`, or `mbm1` perf events fail to boot

If a guest virtual machine is set to report `cmt`, `mbmt`, or `mbm1` perf events, it is unable to boot after the host is upgraded to Red Hat Enterprise Linux 7.5.

To work around this problem, disable this setting by removing lines that contain `event name='cmt'`, `event name='mbmt'`, or `event name='mbm1'` from the `<perf>` section of the domain XML configuration file. (BZ#[1532553](#))

APPENDIX A. COMPONENT VERSIONS

This appendix provides a list of key components and their versions in the Red Hat Enterprise Linux 7.5 release.

Table A.1. Component Versions

Component	Version
kernel	3.10.0-862
kernel-alt	4.14.0-49
QLogic qla2xxx driver	9.00.00.00.07.5-k1
QLogic qla4xxx driver	5.04.00.00.07.02-k0
Emulex lpfc driver	0:11.4.0.4
iSCSI initiator utils (iscsi-initiator-utils)	6.2.0.874-7
DM-Multipath (device-mapper-multipath)	0.4.9-119
LVM (lvm2)	2.02.177-4
qemu-kvm ^[a]	1.5.3-156
qemu-kvm-ma ^[b]	2.10.0-21
<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 Systems. Note that KVM virtualization on IBM POWER9 and IBM z Systems 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# 1274430 , BZ# 1352121 , BZ# 1406351 , BZ# 1458536 , BZ# 1467777 , BZ# 1470169	BZ# 1434335 , BZ# 1445188 , BZ# 1453155 , BZ# 1459946 , BZ# 1464463 , BZ# 1464505 , BZ# 1465600 , BZ# 1476207 , BZ# 1476322 , BZ# 1483681 , BZ# 1498980 , BZ# 1501058 , BZ# 1511462 , BZ# 1517788 , BZ# 1523183 , BZ# 1533571		BZ# 1517383 , BZ# 1544477
Doc-config-command-file-reference	BZ# 1479012			
ModemManager	BZ# 1483051			
NetworkManager	BZ# 1350830 , BZ# 1398925 , BZ# 1436531			
OVMF			BZ# 653382	
OpenIPMI	BZ# 1457805			
adcli		BZ# 1471021		
anaconda	BZ# 1328576 , BZ# 1448459 , BZ# 1450922	BZ# 1452873 , BZ# 1465944 , BZ# 1478970		
ansible			BZ# 1313263	
at		BZ# 1481355		
audit	BZ# 1476406			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
binutils	BZ#1385959, BZ#1406430, BZ#1472955, BZ#1485398	BZ#1465318, BZ#1488889		
checkpolicy	BZ#1494179			
chrony	BZ#1482565			
clevis	BZ#1475406, BZ#1475408, BZ#1478888	BZ#1500975		BZ#1538759
clutter	BZ#1509381			
cockpit	BZ#1470780			
conman	BZ#1435840			
control-center	BZ#1481407			BZ#1543631
corosync			BZ#1413573	
criu			BZ#1400230	
cups	BZ#1434153, BZ#1466497			
curl	BZ#1409208	BZ#1511523		BZ#1510247
custodia			BZ#1403214	
dbus	BZ#1460262, BZ#1480264			BZ#1550582
device-mapper-multipath	BZ#1452210, BZ#1456955	BZ#1459370		
dhcp	BZ#1394727, BZ#1396985			
ding-lib	BZ#1480270			
distribution	BZ#1512020, BZ#1512021			BZ#1062656

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
dnsmasq	BZ# 1188259			
emacs-php-mode	BZ# 1266953			
exiv2		BZ# 1420227		
fence-agents	BZ# 1451776 , BZ# 1476009	BZ# 1519370	BZ# 1476401	
firewalld		BZ# 1462977		
freeipmi	BZ# 1435848			
fwupd	BZ# 1420913			
gcc	BZ# 1535655	BZ# 1468546 , BZ# 1469384 , BZ# 1487434		
gdb		BZ# 1228556 , BZ# 1480498 , BZ# 1493675 , BZ# 1518243		
genwqe-tools		BZ# 1456492		
ghostscript		BZ# 1473337 , BZ# 1479852		
gimp	BZ# 1210840			
gjs		BZ# 1523121		
glibc	BZ# 677316 , BZ# 1375235 , BZ# 1448822 , BZ# 1498925	BZ# 1443236 , BZ# 1504969		
gnome-settings-daemon	BZ# 1481410			
gnome-shell	BZ# 1481381		BZ# 1481395	BZ# 1497303 , BZ# 1511454 , BZ# 1539772 , BZ# 1541021

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
gnome-shell-extensions				BZ# 1544840 , BZ# 1550115
gnome-software				BZ# 1434477
grub2				BZ# 1512493
gssproxy		BZ# 1462974 , BZ# 1488629		
httpd	BZ# 1274890			
hwdata		BZ# 1489281		
ima-evm-utils			BZ# 1384450	
initscripts	BZ# 1357658 , BZ# 1478419	BZ# 1364895 , BZ# 1380496 , BZ# 1395391 , BZ# 1455419		BZ# 1444018
inkscape	BZ# 1480184			
ipa	BZ# 1484683	BZ# 1415162	BZ# 1115294 , BZ# 1298286	BZ# 1478366
ipa-server-docker			BZ# 1405325	
iproute	BZ# 1435647 , BZ# 1456539 , BZ# 1468280			
iptables	BZ# 1402021			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
kernel	BZ#1102454 , BZ#1226051 , BZ#1272615 , BZ#1273769 , BZ#1308630 , BZ#1349668 , BZ#1361287 , BZ#1379551 , BZ#1400689 , BZ#1409365 , BZ#1421164 , BZ#1429710 , BZ#1430637 , BZ#1451916 , BZ#1454745 , BZ#1454965 , BZ#1456687 , BZ#1457561 , BZ#1457572 , BZ#1458278 , BZ#1465223 , BZ#1467288 , BZ#1467335 , BZ#1468286 , BZ#1469857 , BZ#1475409 , BZ#1481303 , BZ#1482253 , BZ#1491226 , BZ#1494476 , BZ#1538911 , BZ#1626526	BZ#947004 , BZ#1317099 , BZ#1373534 , BZ#1383691 , BZ#1432288 , BZ#1438695 , BZ#1442618 , BZ#1442784 , BZ#1445046 , BZ#1446684 , BZ#1448534 , BZ#1450529 , BZ#1457046 , BZ#1460106 , BZ#1460213 , BZ#1460641 , BZ#1462363 , BZ#1465711 , BZ#1467280 , BZ#1467521 , BZ#1467561 , BZ#1469200 , BZ#1469247 , BZ#1472892 , BZ#1476040 , BZ#1476709 , BZ#1479043 , BZ#1506338 , BZ#1507821	BZ#916382 , BZ#1109348 , BZ#1111712 , BZ#1206277 , BZ#1230959 , BZ#1274459 , BZ#1299662 , BZ#1305092 , BZ#1348508 , BZ#1350553 , BZ#1387768 , BZ#1391561 , BZ#1393375 , BZ#1414957 , BZ#1457533 , BZ#1460849	BZ#1303244 , BZ#1367444 , BZ#1470932 , BZ#1484047 , BZ#1486100 , BZ#1509444 , BZ#1528466 , BZ#1535168 , BZ#1539686 , BZ#1540061 , BZ#1540643 , BZ#1548055
kernel-rt	BZ#1401061 , BZ#1462329		BZ#1297061	
kexec-tools	BZ#1431974	BZ#1448861 , BZ#1476219		BZ#1549355
kmod	BZ#1361857			
krb5	BZ#1462982	BZ#1431198 , BZ#1460089		
ksc				BZ#1441455
libdb		BZ#1349779		BZ#1460077

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
libguestfs	BZ#1172425, BZ#1438710, BZ#1448739, BZ#1451665	BZ#1472719, BZ#1506572	BZ#1387213, BZ#1441197, BZ#1477912	
libica	BZ#1376836			
libnftnl			BZ#1332585	
libpfm	BZ#1474999			
libreoffice	BZ#1474303			
libreswan	BZ#1300763, BZ#1457904, BZ#1463062, BZ#1471763, BZ#1475434		BZ#1375750	BZ#1544463
libsmbios	BZ#1463329			
libstoragegmt			BZ#1119909	
libusnic_verbs			BZ#916384	
libva	BZ#1456903			
libvirt	BZ#1289368, BZ#1292451, BZ#1472263		BZ#1283251	BZ#1532553
libvncserver		BZ#1314814		
libyami	BZ#1456906			
linuxptp	BZ#1002657			
logrotate		BZ#1465720		
lorax	BZ#1458937, BZ#1478448	BZ#1341280		
lvm2	BZ#1113681, BZ#1278192			BZ#1337220

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
m17n-db	BZ# 1058510			
mailx		BZ# 1474130		
mod_nss		BZ# 1461580		
mpg123	BZ# 1481753			
mutter	BZ# 1481386			BZ# 1500397 , BZ# 1537529
net-snmp		BZ# 1329338		
netpbm	BZ# 1381122			
nftables	BZ# 1472261	BZ# 1451404		
nmap	BZ# 1460249			
nss	BZ# 1395803 , BZ# 1457789		BZ# 1425514 , BZ# 1431210 , BZ# 1432142	BZ# 1510156 , BZ# 1557015
numpy		BZ# 1167156		
opal-prd	BZ# 1456536			
opencrytpki	BZ# 1456520			
openldap	BZ# 1400578			
opensc	BZ# 1473418			
openscap	BZ# 1505517			BZ# 1556988
openssh	BZ# 1478035	BZ# 1488083 , BZ# 1496808 , BZ# 1517226		BZ# 1493406 , BZ# 1497680
openssl-ibmca	BZ# 1456516			
oprofile	BZ# 1465354			
oscap-anaconda-addon				BZ# 1547609

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
other	BZ# 1432080 , BZ# 1499059 , BZ# 1543995 , BZ# 1578075		BZ# 1062759 , BZ# 1072107 , BZ# 1259547 , BZ# 1464377 , BZ# 1477977	BZ# 1451876 , BZ# 1512940 , BZ# 1515833 , BZ# 1515918 , BZ# 1524193 , BZ# 1532547 , BZ# 1536004
pacemaker	BZ# 1427648 , BZ# 1461976	BZ# 1394418 , BZ# 1489728		
pam		BZ# 1509338		
parted	BZ# 1423357	BZ# 1316239		
pcp	BZ# 1472153			
pcs	BZ# 1367808 , BZ# 1415197	BZ# 1421702 , BZ# 1432283 , BZ# 1508351	BZ# 1433016	
pcsc-lite-ccid	BZ# 1435668			
perl-DBD-MySQL		BZ# 1311646		
perl-DateTime-TimeZone	BZ# 1241818			
perl-HTTP-Daemon		BZ# 1413065		
perl-IO-Socket-SSL	BZ# 1402588			
perl-version		BZ# 1378885		
php	BZ# 1410010			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
pki-core	BZ#1024558, BZ#1400645, BZ#1419761, BZ#1445532, BZ#1446786, BZ#1452347, BZ#1464549, BZ#1469169, BZ#1473452, BZ#1523410, BZ#1523443	BZ#1402280, BZ#1404794, BZ#1446579, BZ#1461217, BZ#1461524, BZ#1465142, BZ#1474658, BZ#1479663, BZ#1484359, BZ#1486225, BZ#1491052, BZ#1498957, BZ#1499054, BZ#1500474, BZ#1506819, BZ#1518096, BZ#1520277, BZ#1532759, BZ#1539125, BZ#1541853		
polycoreutils	BZ#1471809			
python		BZ#1483438		
python-blivet				BZ#1527319
python-urllib3	BZ#1434114			
python-virtualenv	BZ#1461154			
qemu-kvm	BZ#1379822, BZ#1411490	BZ#1455451, BZ#1470244	BZ#1103193	
qemu-kvm-ma	BZ#1400070, BZ#1465503, BZ#1531672			
qgnomeplatform	BZ#1479351			
qt5-qtbase	BZ#1479097			
quota	BZ#1393849			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
rear		BZ# 1388653 , BZ# 1479002 , BZ# 1492177 , BZ# 1506231 , BZ# 1532676		BZ# 1498748 , BZ# 1534646
resource-agents	BZ# 1436189	BZ# 1445628 , BZ# 1457382 , BZ# 1462802		
rhn-client-tools		BZ# 1494389		
rhnlib		BZ# 1503953		
rhnsd		BZ# 1475039 , BZ# 1480306 , BZ# 1489989		
rpcbind				BZ# 1425758
rpm	BZ# 1278924 , BZ# 1406611			
rsync	BZ# 1393543 , BZ# 1432899			
samba	BZ# 1470048			
sane-backends	BZ# 1458903			
sbd	BZ# 1462002 , BZ# 1499864	BZ# 1468580 , BZ# 1525981		
scap-security-guide	BZ# 1404429 , BZ# 1472499			
scap-workbench	BZ# 1479036			BZ# 1533108
selinux-policy	BZ# 1480518 , BZ# 1494172	BZ# 1470735 , BZ# 1472722		
setup	BZ# 1344007	BZ# 1433020		
smartmontools	BZ# 1369731			
sos		BZ# 1183243		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
spice-gtk				BZ# 1540056
squid	BZ# 1452200			
sssd	BZ# 1327705 , BZ# 1400614 , BZ# 1416150 , BZ# 1472255		BZ# 1068725	
strace		BZ# 1466535		
strongimcv			BZ# 755087	
subscription-manager	BZ# 1319927 , BZ# 1329349 , BZ# 1463325 , BZ# 1466453 , BZ# 1499977 , BZ# 1526622	BZ# 1476817 , BZ# 1507158 , BZ# 1519512		
system-config-kdump	BZ# 1384943			
system-config-kickstart		BZ# 1272068		
systemd	BZ# 1384014	BZ# 1455071	BZ# 1284974	
systemtap	BZ# 1473722			
tang	BZ# 1478895			
tboot	BZ# 1457529			
tcpdump	BZ# 1464390 , BZ# 1490842			
tftp	BZ# 1328827			
tpm2-abrmd	BZ# 1492466			
tpm2-tss	BZ# 1463097			
tss2			BZ# 1384452	
tuned	BZ# 1467576			BZ# 1385838

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
unbound	BZ# 1251440			
usbguard			BZ# 1480100	
valgrind	BZ# 1473725			
vdo	BZ# 1480047			
vim	BZ# 1267826 , BZ# 1319760			
virt-manager	BZ# 1472271			
virt-what	BZ# 1476878			
virt-who	BZ# 1408556 , BZ# 1436617	BZ# 1389729 , BZ# 1461417 , BZ# 1485865		
wayland			BZ# 1481411	
webkitgtk4	BZ# 1476707			
xorg-x11-drivers				BZ# 1516970
xorg-x11-drv-intel				BZ# 1545550
xorg-x11-server				BZ# 1478625 , BZ# 1499129
yum	BZ# 1432319	BZ# 1458841		
yum-utils	BZ# 1437636 , BZ# 1470647	BZ# 1428210 , BZ# 1455318		

APPENDIX C. REVISION HISTORY

Revision 0.1-7	Thu Sep 13 2018	Lenka Špačková
Moved CephFS from Technology Previews to fully supported features (File Systems).		
Revision 0.1-6	Tue Aug 28 2018	Lenka Špačková
Fixed a known issue related to performance loss caused by certain security patches (Kernel).		
Revision 0.1-5	Tue Jul 31 2018	Lenka Špačková
Added a note regarding a change in behavior of the reposync command to New Features (System and Subscription Management).		
Revision 0.1-4	Tue Jul 17 2018	Lenka Špačková
Added a known issue related to LVM handling of incomplete volume groups (Storage).		
Revision 0.1-3	Wed Jul 04 2018	Lenka Špačková
Fixed an option name in the DNS stub resolver improvements description (Networking).		
Revision 0.1-2	Wed Jun 13 2018	Lenka Špačková
Fixed cross-reference links.		
Revision 0.1-1	Fri May 18 2018	Lenka Špačková
Removed a duplicate description. Updated GNOME Shell rebase description.		
Revision 0.1-0	Tue May 15 2018	Lenka Špačková
Shenandoah garbage collector moved to fully supported features (Compiler and Tools). Added a description of clufter rebase and of support for Sybase ASE failover to features (Clustering). Added a known issue related to read-only system configuration (Installation and Booting). Added several bug fix descriptions (Clustering, Storage). Expanded wayland Technology Preview description (Desktop).		
Revision 0.0-9	Tue Apr 24 2018	Lenka Špačková
Moved the tpm2-* packages from Technology Previews to fully supported features (Hardware Enablement). Added a new OpenSC feature related to CAC Alternate tokens (Security). Added a known issue related to NSS and smart cards (Security).		
Revision 0.0-8	Tue Apr 17 2018	Lenka Špačková
Updated a recommendation related to the sslwrap() deprecation. Added a PTP device addition note (Virtualization).		
Revision 0.0-7	Fri Apr 13 2018	Lenka Špačková
Updated a link to Intel® Omni-Path Architecture documentation.		
Revision 0.0-6	Tue Apr 10 2018	Lenka Špačková
Release of the Red Hat Enterprise Linux 7.5 Release Notes.		
Revision 0.0-1	Wed Jan 24 2018	Lenka Špačková
Release of the Red Hat Enterprise Linux 7.5 Beta Release Notes.		