



Red Hat Enterprise Linux 7

7.4 Release Notes

Release Notes for Red Hat Enterprise Linux 7.4

Red Hat Enterprise Linux 7 7.4 Release Notes

Release Notes for Red Hat Enterprise Linux 7.4

Red Hat Customer Content Services

rhel-notes@redhat.com

Legal Notice

Copyright © 2017-2018 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 7.4 and document known problems in this release, as well as notable bug fixes, Technology Previews, deprecated functionality, and other details.

Table of Contents

PREFACE	16
CHAPTER 1. OVERVIEW	17
Security	17
Identity Management	17
Networking	17
Kernel	18
Storage and File Systems	18
Tools	18
High Availability	18
Virtualization	18
Management and Automation	18
Red Hat Insights	19
Red Hat Customer Portal Labs	19
CHAPTER 2. ARCHITECTURES	20
CHAPTER 3. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	21
UPDATED /PROC/SYS/KERNEL ENTRIES	21
UPDATED /PROC/SYS/USER ENTRIES	21
KERNEL PARAMETERS	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
cloud-init moved to the Base channel	25
CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY	26
SSSD in a container now fully supported	26
Identity Management now supports FIPS	26
SSSD supports obtaining a Kerberos ticket when users authenticate with a smart card	26
SSSD enables logging in to different user accounts with the same smart card certificate	26
IdM web UI enables smart card login	26
New packages: keycloak-httpd-client-install	27
New Kerberos credential cache type: KCM	27
AD users can log in to the web UI to access their self-service page	27
SSSD enables configuring an AD subdomain in the SSSD server mode	27
SSSD supports user and group lookups and authentication with short names in AD environments	28
SSSD supports user and group resolution, authentication, and authorization in setups without UIDs or SIDs	28
SSSD introduces the ssctl user-checks command, which checks basic SSSD functionality in a single operation	29
Support for secrets as a service	29
IdM enables semi-automatic upgrades of the IdM DNS records on an external DNS server	29
IdM now generates SHA-256 certificate and public key fingerprints	29
IdM supports flexible mapping mechanisms for linking smart card certificates to user accounts	29
New user-space tools enable a more convenient LMDB debugging	30
openldap rebased to version 2.4.44	30
Improved security of DNS lookups and robustness of service principal lookups in Identity Management	30
samba rebased to version 4.6.2	30
authconfig can enable SSSD to authenticate users with smart cards	31
authconfig can now enable account locking	31
Improved performance of the IdM server	31

The default session expiration period in the IdM web UI has changed	31
The dbmon.sh script now uses instance names to connect to Directory Server instances	32
Directory Server now uses the SSHA_512 password storage scheme as default	32
Directory Server now uses the tcmmalloc memory allocator	32
Directory Server now uses the nunc-stans framework	32
Improved performance of the Directory Server memberOf plug-in	32
Directory Server now logs severity levels in the error log file	32
Directory Server now supports the PBKDF2_SHA256 password storage scheme	33
Improved auto-tuning support in Directory Server	33
New PKI configuration parameter allows control of the TCP keepalive option	33
PKI Server now creates PKCS #12 files using strong encryption	33
CC-compliant algorithms available for encryption operations	33
New options to allow configuring visibility of menu items in the TPS interface	33
Added a profile component to copy certificate Subject Common Name to the Subject Alternative Name extension	34
New option to remove LDAP entries before LDIF import	34
Certificate System now supports externally authenticated users	34
Certificate System now supports enabling and disabling certificate and CRL publishing	34
The searchBase configuration option has been added to the DirAcIAuthz PKI Server plug-in	35
For better performance, Certificate System now supports ephemeral	35
Section headers in PKI deployment configuration file are no longer case sensitive	35
Certificate System now supports installing a CA using HSM on FIPS-enabled Red Hat Enterprise Linux	35
CMC requests now use a random IV for AES and 3DES encryption	35
CHAPTER 6. CLUSTERING	36
clutter rebased to version 0.76.0 and fully supported	36
Support for quorum devices in a Pacemaker cluster	36
Support for Booth cluster ticket manager	37
Support added for using shared storage with the SBD daemon	37
Full support for CTDB resource agent	37
The High Availability and Resilient Storage Add-Ons are now available for IBM POWER, little endian	37
pcs now provides the ability to set up a cluster with encrypted corosync communication	37
New commands for supporting and removing remote and guest nodes	37
Ability to configure pcsd bind addresses	37
New option to the pcs resource unmanage command to disable monitor operations	38
Support for regular expressions in pcs command line when configuring location constraints	38
Specifying nodes in fencing topology by a regular expression or a node attribute and its value	38
Support for Oracle 11g for the resource agents Oracle and OraLsnr	38
Support for using SBD with shared storage	38
Support for NodeUtilization resource agent	38
CHAPTER 7. COMPILER AND TOOLS	40
pcp rebased to version 3.11.8	40
systemtap rebased to version 3.1	40
valgrind rebased to version 3.12	40
New package: unitsofmeasurement	41
SSL/TLS certificate verification for HTTP clients is now enabled by default in the Python standard library	41
Support for %gemspec_add_dep and %gemspec_remove_dep has been added	41
ipmitool rebased to version 1.8.18	41
lshw updated for the little-endian variant of IBM Power	41
perf now supports uncore events on Intel Xeon v5	41
dmidecode updated	41
iSCSI now supports configuring the ALUA operation by using targetcli	41

jansson rebased to version 2.10	42
A new compatibility environmental variable for egrep and fgrep	42
lastcomm now supports the --pid option	42
New package: perl-Perl4-CoreLibs	42
tar now follows symlinks to directories when extracting from the archive	42
The IO::Socket::SSL Perl module now supports restricting of TLS version	42
The Net::SSLeay Perl module now supports restricting of TLS version	42
wget now supports specification of the TLS protocol version	42
tcpdump rebased to version 4.9.0	43
The option to set capture direction for tcpdump changed from -P to -Q	43
OpenJDK now supports SystemTap on the 64-bit ARM architecture	43
sos rebased to version 3.4	43
targetd rebased to version 0.8.6	43
shim rebased to version 12-1	44
rubygem-abrt rebased to version 0.3.0	44
New package: http-parser	44
Intel and IBM POWER transactional memory support for all default POSIX mutexes	44
glibc now supports group merging	44
glibc now supports optimized string comparison functions on The IBM POWER9 architecture	44
Improved performance for dynamically loaded libraries using the Intel SSE, AVX and AVX512 features	44
elfutils rebased to version 0.168	45
bison rebased to version 3.0.4	45
The system default CA bundle has been set as default in the compiled-in default setting or configuration in Mutt	45
objdump mixed listing speed up	45
ethtool support for human readable output from the fjes driver	46
ecj rebased to version 4.5.2	46
rhino rebased to version 1.7R5	46
scap-security-guide and oscap-docker now support containers	46
CHAPTER 8. DESKTOP	47
GNOME rebased to version 3.22.3	47
The xorg-x11-drv-libinput driver has been added to the X.Org input drivers	47
Change of default driver for some Intel and nVidia Hardware	47
dconf-editor is now provided by a separate package	47
CHAPTER 9. FILE SYSTEMS	48
SELinux security labels are now supported on the OverlayFS file system	48
NFSv4.1 server is now fully supported	48
autofs now supports the browse options of amd format maps	48
To make searching logs easier, autofs now provides identifiers of mount request log entries	48
GFS2 on IBM z Systems is now supported in SSI environments	48
gfs2-utils rebased to version 3.1.10	48
FUSE now supports SEEK_HOLE and SEEK_DATA in lseek calls	49
NFS server now supports limited copy-offload	49
SELinux is supported for use with GFS2 file systems	49
NFSv4.1 client and server now support Kerberos authentication	49
rpc.idmapd now supports obtaining NFSv4 ID Domains from DNS	49
NFSv4.1 is now the default NFS mount protocol	49
Setting nfs-utils configuration options has been centralized in nfs.conf	50
Locking performance for NFSv4.1 mounts has been improved for certain workloads	50
CHAPTER 10. HARDWARE ENABLEMENT	51
Hardware utility tools now correctly identify recently released hardware	51

New Wacom driver introduced in 7.4 to support upcoming tablets	51
Wacom kernel driver now supports ThinkPad X1 Yoga touch screen	51
The touch functionality has been added to the Wacom Cintiq 27 QHDT tablets	51
AMDGPU now supports the Southern Islands, Sea Islands, Volcanic Islands and Arctic Islands chipsets	51
Support added for the AMD mobile graphics	51
Netronome NFP devices are supported	51
nvme-cli rebased to version 1.3	51
The queued spinlocks have been implemented into the Linux kernel	51
rapl now supports Intel Xeon v2 servers	52
Further support for Intel Platform Controller Hub [PCH] devices	52
Included genwqe-tools to enable use of hardware accelerated zLib on IBM Power and s390x	52
librtas rebased to version 2.0.1	52
The NFP driver	52
Enable latest NVIDIA cards in Nouveau	52
Support for Wacom ExpressKey Remote	52
Wacom Cintiq 27 QHD now supports ExpressKey Remote	52
CHAPTER 11. INSTALLATION AND BOOTING	53
Anaconda enables users to set RAID chunk size	53
Anaconda text mode now supports IPoIB interfaces	53
inst.debug enables a more convenient debugging of Anaconda installation issues	53
Kickstart installation failure automatically triggers %onerror scripts	53
Anaconda can now wait for network to become available before starting the installation	53
Multiple network locations of stage2 or Kickstart files can be specified to prevent installation failure	53
autopart --nohome in a kickstart file disables the creation of /home/ in automatic partitioning	53
Loading driver disks from hard disk drives and USBs enabled	54
Changes in automatic partitioning behavior for LVM thin pools	54
32-bit boot loaders can now boot 64-bit kernels on UEFI	54
Lorax can now ignore SSL errors	54
shim-signed rebased to version 12	55
gnu-efi rebased to version 3.0.5.-9	55
Backward compatibility enabled for killproc() and status()	55
DHCP_FQDN allows specifying a fully qualified domain name of the system	55
You can now create thin logical volume snapshots during the installation process	55
CHAPTER 12. KERNEL	56
The NVMe driver rebased to kernel version 4.10	56
crash rebased to version 7.1.9	56
crash now analyzes vmcore dumps for IBM Power ISA 3.0	56
crash updated for IBM Power and for the little-endian variant of IBM Power	56
memkind updated to version 1.3.0	56
Jitter Entropy RNG added to the kernel	56
/dev/random now shows notifications and warnings for the urandom pool initialization	57
fjes updated to version 1.2	57
Full support for user name spaces	57
makedumpfile updated to version 1.6.1	57
Intel Cache Allocation Technology is supported	57
qat updated to the latest upstream version	57
Addition of intel-cmt-cat package	57
i40e now supports trusted and untrusted VFs	57
Kernel support for OVS 802.1ad (QinQ)	57
Live post-copy migration support for shared memory and hugetlbfs	58
New package: dbxtool	58

mlx5 now supports SRIOV-trusted VFs	58
rwsem performance updates from the 4.9 kernel backported	58
getrandom added to the Linux kernel	58
A new status line, Umask, has been included in /proc/<PID>/status	58
Intel® Omni-Path Architecture (OPA) host software	58
The XTS-AES key verification now meets the FIPS 140-2 requirements	59
mlx5 is now supported on IBM z Systems	59
The perf tool now supports processor cache-line contention detection	59
SCSI-MQ support in the lpfc driver	59
CHAPTER 13. REAL-TIME KERNEL	60
About Red Hat Enterprise Linux for Real Time Kernel	60
kernel-rt rebased	60
CHAPTER 14. NETWORKING	61
NetworkManager rebased to version 1.8	61
NetworkManager now supports additional features for routes	61
NetworkManager now better handles devices state	61
NetworkManager now supports MACsec (IEEE 802.1AE)	61
NetworkManager now supports changing and enforcing 802-3 link properties	61
NetworkManager now supports ordering bond slaves based on device names	61
NetworkManager now supports VFs for SR-IOV devices	62
Kernel GRE rebased to version 4.8	62
dnsmasq rebased to version 2.76	62
BIND changes the way it handles URI resource records, impacting also URI backward compatibility	63
A DHCP client hook example added for DDNS for Microsoft Azure cloud	63
dhcp_release6 now releases IPv6 addresses	63
Sendmail now supports ECDHE	63
telnet now supports the -6 option	63
Adjustable TTL limit for caching negative DNS responses in Unbound	63
The scalability of UDP sockets has been improved	63
IP now supports IP_BIND_ADDRESS_NO_PORT in the kernel	64
IPVS Source Hash scheduling now supports L4 hashing and SH fallback	64
iproute now supports changing bridge port options	64
New options of Sockets API Extensions for SCTP (RFC 6458) implemented	64
ss now supports SCTP sockets list	64
wpa_supplicant rebased to version 2.6	64
Linux kernel now contains the switchdev infrastructure and mlxsw	64
The Linux bridge code rebased to version 4.9	65
bind-dyndb-ldap rebased to version 11.1	65
DynDB API from the upstream version 9.11.0 of BIND added to Red Hat Enterprise Linux	66
tboot rebased to version 1.9.5	66
Packages related to rdma consolidated by rebase into rdma-core version 13	66
OVN IP address management support added for static MAC addresses	67
Enhanced networked reliability on multihomed hosts	67
Offloading of GENEVE, VXLAN, and GRE tunnels is now supported	67
LCO for tunnel traffic is now supported	67
Improved tunnel performance on NICs	67
NPT is now supported in the kernel	67
DNS configuration is now supported through the D-Bus API	68
PPP support is now moved into a separate package	68
The tc utility now supports flower	68
Fix to the CRC32c value computation in SCTP forwarding path	68

New packages: iperf3	68
Installation of OVN now supports easily-configurable firewalld rules	68
netlink now supports bridge master attributes	68
CHAPTER 15. SECURITY	69
New packages: tang, clevis, jose, luksmeta	69
New package: usbguard	69
openssh rebased to version 7.4	69
audit rebased to version 2.7.6	70
opensc rebased to version 0.16.0	70
openssl rebased to version 1.0.2k	71
openssl-ibmca rebased to version 1.3.0	71
OpenSCAP 1.2 is NIST-certified	71
libreswan rebased to version 3.20	71
Audit now supports filtering based on session ID	72
libseccomp now supports IBM Power architectures	72
AUDIT_KERN_MODULE now records module loading	72
OpenSSH now uses SHA-2 for public key signatures	72
firewalld now supports additional IP sets	72
firewalld now supports actions on ICMP types in rich rules	73
firewalld now supports disabled automatic helper assignment	73
nss and nss-util now use SHA-256 by default	73
Audit filter exclude rules now contain additional fields	73
PROCTITLE now provides the full command in Audit events	73
nss-softoken rebased to version 3.28.3	73
libica rebased to version 3.0.2	74
opencryptoki rebased to version 3.6.2	74
AUDIT_NETFILTER_PKT events are now normalized	74
p11tool now supports writing objects by specifying a stored ID	74
new package: nss-pem	74
pmrfc3164 replaces pmrfc3164sd in rsyslog	74
libreswan now supports right=%opportunisticgroup	75
ca-certificates now meet Mozilla Firefox 52.2 ESR requirements	75
nss now meets Mozilla Firefox 52.2 ESR requirements for certificates	75
scap-security-guide rebased to version 0.1.33	75
CHAPTER 16. SERVERS AND SERVICES	77
chrony rebased to version 3.1	77
linuxptp rebased to version 1.8	77
tuned rebased to version 2.8.0	77
logrotate now uses /var/lib/logrotate/logrotate.status as the default state file	77
rsyslog rebased to version 8.24.0	78
New cache configuration options for mod_nss	78
Database and prefix options have been removed from nss_pcache	79
New package: libfastjson	79
tuned now supports initrd overlays	79
openwsman now supports disabling of particular SSL protocols	79
rear rebased to version 2.0	79
python-tornado rebased to version 4.2.1	80
CHAPTER 17. STORAGE	81
Support added in LVM for RAID level takeover	81
LVM now supports RAID reshaping	81
Device Mapper linear devices now support DAX	81

libstoragemgmt rebased to version 1.4.0	81
mpt3sas updated to version 15.100.00.00	81
The lpfc_no_hba_reset module parameter for the lpfc driver is now available	81
LVM now detects Veritas Dynamic Multi-Pathing systems and no longer accesses the underlying device paths directly	82
The libnvdimm kernel subsystem now supports PMEM subdivision	82
Warning messages when multipathd is not running	82
c library interface added to multipathd to give structured output	82
New remove retries multipath configuration value	82
New multipathd reset multipaths stats commands	82
New disable_changed_wwids mulitpath configuration parameter	82
Updated built-in configuration for HPE 3PAR array	83
Added built-in configuration for NFINIDAT InfiniBox.* devices	83
device-mapper-multipath now supports the max_sectors_kb configuration parameter	83
New detect_checker multipath configuration parameter	83
Multipath now has a built-in default configuration for Nimble Storage devices	83
LVM supports reducing the size of a RAID logical volume	83
iprutils rebased to version 2.4.14	83
mdadm rebased to version 4.0	84
LVM extends the size of a thin pool logical volume when a thin pool fills over 50 percent	84
LVM now supports dm-cache metadata version 2	84
Support for DIF/DIX (T10 PI) on specified hardware	84
The dmstats facility can now track the statistics for files that change	85
Support for thin snapshots of cached logical volumes	85
New package: nvmetcli	85
Device DAX is now available for NVDIMM devices	85
CHAPTER 18. SYSTEM AND SUBSCRIPTION MANAGEMENT	86
New payload_gpgcheck option added to yum	86
A no-proxy configuration is available for virt-who	86
virt-who respects independent interval settings	86
Password options added to virt-who-password	86
Regular expressions and wildcards can be used in some virt-who configuration parameters	86
virt-who configuration files are easier to manage	86
CHAPTER 19. VIRTUALIZATION	87
ENA drivers for Amazon Web Services	87
Synthetic Hyper-V FC adapters are supported by the storvsc driver	87
Parent HBA can be defined by a WWNN/WWPN pair	87
libvirt rebased to version 3.2.0	87
KVM now supports MCE	87
Added support for rx batching on tun/tap devices	87
libguestfs rebased to version 1.36.3	87
Improved virt-v2v installation of QXL drivers	88
virt-v2v can export disk images to qcow2 format 1.1	88
Additional virt tools can work on LUKS whole-disk encrypted guests	88
Tab completion for all libguestfs commands	88
Resized disks can be written directly to a remote location	88
User namespace is now fully supported	88
Driver added for devices that connect over a PCI Express bus in guest virtual machine under Hyper-V	88
CHAPTER 20. ATOMIC HOST AND CONTAINERS	89
Red Hat Enterprise Linux Atomic Host	89

CHAPTER 21. RED HAT SOFTWARE COLLECTIONS	90
PART II. NOTABLE BUG FIXES	91
CHAPTER 22. GENERAL UPDATES	92
Addition of CtrlAltDelBurstAction for Systemd	92
cgred can now resolve rules concerning NSS users and groups	92
CHAPTER 23. AUTHENTICATION AND INTEROPERABILITY	93
yum no longer reports package conflicts after installing ipa-client	93
In FIPS mode, the slapd_pk11_getInternalKeySlot() function is now used to retrieve the key slot for a token	93
Certificate System no longer fails to install with a Thales HSM on systems in FIPS mode	93
The dependency list for pkispawn now correctly includes openssl	93
Error messages from the PKI Server profile framework are now passed through to the client	93
Certificate System does not start a Lightweight CA key replication during installation	93
PKI Server now correctly compares subject DNs during startup	93
KRA installation no longer fails when connecting to an intermediate CA with an incomplete certificate chain	94
The startTime field in certificate profiles now uses long integer format	94
Subordinate CA installation no longer fails due with a PKCS#11 token is not logged in error	94
The pkispawn script now correctly sets the ECC key sizes	94
CA clone installation in FIPS mode no longer fails	94
PKI Server no longer fails to start when an entryUSN attribute contains a value larger than 32-bit	94
Tomcat now works with IPv6 by default	94
pkispawn no longer generates invalid NSS database passwords	95
Certificate retrieval no longer fails when adding a user certificate with the --serial option	95
CA web interface no longer shows a blank certificate request page if there is only one entry	95
Installing PKI Server in a container environment no longer displays a warning	95
Re-enrolling a token using a G&D smart card no longer fails	95
PKI Server provides more detailed information about certificate validation errors on startup	95
PKI Server no longer fails to re-initialize the LDAPProfileSubsystem profile	95
Extracting private keys generated on an HSM no longer fails	95
pkispawn no longer generates passwords consisting only of digits	96
CA certificates are now imported with correct trust flags	96
Generating a symmetric key no longer fails when using the --usage verify option	96
Subsequent PKI installation no longer fails	96
Two-step subordinate CA installation in FIPS mode no longer fails	96
The audit log no longer records success when a certificate request was rejected or canceled	96
PKI subsystems which failed self tests are now automatically re-enabled on startup	96
CERT_REQUEST_PROCESSED audit log entries now include certificate serial number instead of encoded data	97
Updating the LDAPProfileSubsystem profile now supports removing attributes	97
CHAPTER 24. CLUSTERING	98
Pacemaker Remote may shut down, even if its connection to the cluster is unmanaged	98
pcs now validates the name and the host of a remote and guest node	98
New syntax of master option in pcs resource create command allows correct creation of meta attributes	98
CHAPTER 25. COMPILER AND TOOLS	99
The PCRE library now correctly recognizes non-ASCII printable characters as required by Unicode	99
Applications using Bundler to manage dependencies can now properly load the JSON library	99
Git can now be used with HTTP or HTTPS and SSO	99
rescan-scsi-bus.sh --luns=1 now scans only LUNs numbered with 1	99
ps no longer removes prefixes from wait channel names	99
tcsh no longer becomes unresponsive when the .history file is located on a network file system	99

fcoeadm --target no longer causes fcoeadm to crash	100
tar option --directory no longer ignored	100
tar options --xattrs-exclude and --xattrs-include no longer ignored	100
tar now restores incremental backup correctly	100
The perl-homedir profile scripts now support csh	100
getaddrinfo no longer accessing uninitialised data	100
Additional security checks performed in malloc implementation in glibc	100
chrpath rebased to version 0.16	101
Updated translations for the system-config-language package	101
Mutt no longer send emails with an incomplete From header when a host name lacks the domain part	101
strace displays correctly the O_TMPFILE flag and mode for open() function	101
ld no longer enters an infinite loop when linking large programs	101
gold warning messages for cross object references to hidden symbols fixed	101
OProfile default event on Intel Xeon® C3xxx Processors with Denverton SOC fixed	101
CHAPTER 26. DESKTOP	102
Empathy now can validate certificate chains for Google Talk	102
CHAPTER 27. FILE SYSTEMS	103
Setting the retry timeout can now prevent autofs from starting without mounts from SSSD	103
The autofs package now contains the README.autofs-schema file and an updated schema	103
automount no longer needs to be restarted to access maps stored on the NIS server	103
Checking local mount availability with autofs no longer leads to a lengthy timeout before failing	103
The journal is marked as idle when mounting a GFS2 file system as read-only	103
The id command no longer shows incorrect UIDs and GIDs	104
Labeled NFS is now turned off by default	104
autofs mounts no longer enter an infinite loop after reaching a shutdown state	104
autofs is now more reliable when handling namespaces	104
CHAPTER 28. INSTALLATION AND BOOTING	105
Automatic partitioning now works when installing on a single FBA DASD on IBM z Series	105
Activation of bridge configured in Kickstart no longer fails when Kickstart proceeds from the disk	105
Anaconda now correctly allows creating users without passwords	105
Minimal installation no longer installs open-vm-tools-desktop and dependencies	105
Anaconda no longer generates invalid Kickstart files	105
Anaconda no longer fails to identify RAID arrays specified by name	105
Kickstart no longer accepts passwords that are too short	106
Initial Setup now correctly opens in a graphical interface over SSH on IBM z Systems	106
Extra time is no longer needed for installation when geolocation services are enabled	106
The ifup-aliases script now sends gratuitous ARP updates when adding new IP addresses	106
The netconsole utility now launches correctly	106
rc.debug kernel allows easier debugging of initscripts	106
The system no longer fails to terminate with /usr on iSCSI or NFS	106
rhel-autorelabel no longer corrupts the filesystem	106
The rpmbuild command now correctly processes Perl requires	107
Installer now correctly recognizes BIOS RAID devices when using ignoredisk in Kickstart	107
Single quotes now work for values in the ifcfg-* files	107
rhel-import-state no longer changes access permissions for /dev/shm/, allowing the system to boot correctly	107
Backward compatibility enabled for Red Hat Enterprise Linux 6 initscripts	107
initscripts now specifies /etc/rwtab and /etc/statetab as configuration files	107
The ifup script no longer slows down NetworkManager	107
Gnome Initial Setup can now be disabled by the firstboot --disable command in kickstart	108
Setting NM_CONTROLLED now works correctly across all the ifcfg-* files	108

The dhclient command no longer incorrectly uses localhost when hostname is not set	108
The initscripts utility now handles LVM2 correctly	108
The service network stop command no longer attempts to stop services which are already stopped	108
ifdown on a loopback device now works correctly	108
Scripts in initscripts handle static IPv6 address assignment more robustly	108
Deselecting an add-on option in Software Selection no longer requires a double-click	108
The target system hostname can be configured via installer boot options in Kickstart installations	109
Anaconda no longer asks for Installation Source verification after network configuration	109
Disks using the OEMDRV label are now correctly ignored during automatic installation	109
CHAPTER 29. KERNEL	110
RAID 4 and RAID 10 creation and activation fully supported	110
kdump now works with legacy type 12 NVDIMMs	110
Creating a file that inherits ACLs no longer loses mask	110
CHAPTER 30. REAL-TIME KERNEL	111
Removing USB no longer causes a might_sleep() warning on MRG Realtime kernel	111
CHAPTER 31. NETWORKING	112
SNMP response is no longer timed out	112
ICMP redirects no longer cause kernel to crash	112
The net.ipv4.ip_nonlocal_bind kernel parameter is set in name spaces	112
The netfilter REJECT rule now works on SCTP packets	112
NetworkManager no longer duplicates a connection with already-set DHCP_HOSTNAME	112
Improved SCTP congestion_window management	112
Value of DCTCP alpha now drops to 0 and cwnd remains at values more than 137	112
ss now displays correctly cwnd	113
Value of cwnd no longer increases using DCTCP	113
Negated range matches have been fixed	113
The nmcli connection show command now displays the correct output for both empty and NULL values	113
snmpd no longer rejects large packets from AgentX subagents	113
Macvlan can now be unregistered correctly	113
CHAPTER 32. SECURITY	114
Configurations that depend on chrooting in user-non-searchable paths now work properly	114
firewalld now supports all ICMP types	114
docker.pp replaced with container.pp in selinux-policy	114
Recently-added kernel classes and permission defined in selinux-policy	114
nss now properly handles PKCS#12 files	114
OpenSCAP now produces only useful messages and warnings	114
AIDE now logs in the syslog format	114
Installations with the OpenSCAP security-hardening profile now proceed	114
OpenSCAP and SSG are now able to scan RHV-H systems correctly	115
OpenSCAP now handles also uncompressed XML files in a CVE OVAL feed	115
CHAPTER 33. SERVERS AND SERVICES	116
rear now correctly preserves Linux capabilities	116
sblim-cmpi-fsvol no longer shows file systems mounted with DM as disabled	116
SPNEGO in Cyrus SASL is now compatible with Microsoft Windows	116
Data are no longer lost when the MariaDB init script fails	116
ypbind no longer starts before access to the network is guaranteed	116
Remote users' account settings are no longer reverted to default settings on restart due to ypbind	116
ypasswd no longer crashes due to Network Information System security features used	117
Evince now displays PostScript files again	117

db_verify no longer causes libdb to run out of free mutexes	117
ghostscript no longer becomes unresponsive in some situations	117
Converting postscript to PDF no longer causes ps2pdf to terminate unexpectedly	117
sapconf now works correctly with higher kernel.shmall and kernel.shmmax values	117
CHAPTER 34. STORAGE	118
lvconvert --repair now works properly on cache logical volumes	118
LVM2 library incompatibilities no longer cause device monitoring to fail and be lost during an upgrade	118
be2iscsi driver errors no longer cause the system to become unresponsive	118
Interaction problems no longer occur with the lvmetad daemon when mirror segment type is used	118
The multipathd daemon no longer shows incorrect error messages for blacklisted devices	118
Multipath now flags device reloads when there are no usable paths	118
Read requests sent after failed writes will always return the same data on multipath devices	118
When a path device in a multipath device switches to read-only, the multipath device will be reloaded read-only	119
Users no longer get potentially confusing stale data for multipath devices that are not being checked	119
The multipathd daemon no longer hangs as a result of running the prioritizer on failed paths	119
New RAID4 volumes, and existing RAID4 or RAID10 logical volumes after a system upgrade are now correctly activated	119
LVM tools no longer crash due to an incorrect status of PVs	119
CHAPTER 35. SYSTEM AND SUBSCRIPTION MANAGEMENT	120
Undercloud no longer fails on a system with no configured repositories	120
the yum commands provided by the yum-plugin-verify now set the exit status to 1 if any mismatches are found	120
CHAPTER 36. VIRTUALIZATION	121
SeaBIOS recognizes SCSI devices with a non-zero LUN	121
The libguestfs tools now correctly handle guests where /usr/ is not on the same partition as root	121
virt-v2v can convert Windows guests with corrupted or damaged Windows registries	121
Converting Windows guests with non-system dynamic disks using virt-v2v now works correctly	121
Guests can be converted to Glance images, regardless of the Glance client version	121
Red Hat Enterprise Linux 6.2 - 6.5 guest virtual machines can now be converted using virt-v2v	121
Btrfs entries in /etc/fstab are now parsed correctly by libguestfs	121
libguestfs can now correctly open libvirt domain disks that require authentication	122
Converted Windows UEFI guests boot properly	122
The virt-v2v utility now ignores proxy environment variables consistently	122
virt-v2v only copies rhev-apt.exe and rhsrvany.exe when needed	122
Guests with VLAN over a bonded interface no longer stop passing traffic after a failover	122
virt-v2v imports OVAs that do not have the <ovf:Name> attribute	122
PART III. TECHNOLOGY PREVIEWS	123
CHAPTER 37. GENERAL UPDATES	124
The systemd-importd VM and container image import and export service	124
CHAPTER 38. AUTHENTICATION AND INTEROPERABILITY	125
Use of AD and LDAP sudo providers	125
DNSSEC available as Technology Preview in IdM	125
Identity Management JSON-RPC API available as Technology Preview	125
The Custodia secrets service provider is now available	125
Containerized Identity Management server available as Technology Preview	126
CHAPTER 39. CLUSTERING	127
The pcs tool now manages bundle resources in Pacemaker	127

CHAPTER 40. COMPILER AND TOOLS	128
Shenandoah garbage collector	128
CHAPTER 41. FILE SYSTEMS	129
The CephFS kernel client is now available	129
ext4 and XFS file systems now support DAX	129
pNFS and Block Layout Support	129
OverlayFS	129
pNFS SCSI layouts client and server support is now provided	130
Btrfs file system	130
CHAPTER 42. HARDWARE ENABLEMENT	131
Trusted Computing Group TPM 2.0 System API library and management utilities available	131
New package: tss2	131
LSI Syncro CS HA-DAS adapters	131
CHAPTER 43. INSTALLATION AND BOOTING	132
Multi-threaded xz compression in rpm-build	132
CHAPTER 44. KERNEL	133
Heterogeneous memory management included as a Technology Preview	133
criu rebased to version 2.12	133
kexec as a Technology Preview	133
kexec fast reboot as a Technology Preview	133
Unprivileged access to name spaces can be enabled as a Technology Preview	133
KASLR as a Technology Preview	133
Updated NFSv4 pNFS clients with flexible file layout	134
CUIR enhanced scope detection	134
SCSI-MQ as a Technology Preview in the qla2xxx driver	134
CHAPTER 45. REAL-TIME KERNEL	135
New scheduler class: SCHED_DEADLINE	135
CHAPTER 46. NETWORKING	136
Cisco usNIC driver	136
Cisco VIC kernel driver	136
Trusted Network Connect	136
SR-IOV functionality in the qlcnict driver	136
The libnftnl and nftables packages	136
The flower classifier with off-loading support	136
CHAPTER 47. RED HAT ENTERPRISE LINUX SYSTEM ROLES POWERED BY ANSIBLE	137
New packages: ansible	137
CHAPTER 48. SECURITY	138
The tang-nagios and clevis-udisk2 subpackages available as a Technology Preview	138
USBGuard is now available for IBM Power as a Technology Preview	138
CHAPTER 49. STORAGE	139
Multi-queue I/O scheduling for SCSI	139
Targetd plug-in from the libStorageMgmt API	139
Support for Data Integrity Field/Data Integrity Extension (DIF/DIX)	139
CHAPTER 50. VIRTUALIZATION	140
USB 3.0 support for KVM guests	140
Select Intel network adapters now support SR-IOV as a guest on Hyper-V	140

No-IOMMU mode for VFIO drivers	140
The ibmvnic Device Driver has been added	140
virt-v2v can now use vmx configuration files to convert VMware guests	140
virt-v2v can convert Debian and Ubuntu guests	140
Virtio devices can now use vIOMMU	141
Open Virtual Machine Firmware	141
PART IV. DEVICE DRIVERS	142
CHAPTER 51. NEW DRIVERS	143
Storage Drivers	143
Network Drivers	143
Graphics Drivers and Miscellaneous Drivers	144
CHAPTER 52. UPDATED DRIVERS	145
Storage Driver Updates	145
Network Driver Updates	145
Graphics Driver and Miscellaneous Driver Updates	146
CHAPTER 53. DEPRECATED FUNCTIONALITY	147
Deprecated packages related to Identity Management	147
Deprecated Insecure Algorithms and Protocols	147
Legacy CA certificates removed from the ca-certificates package	150
coolkey replaced with opensc	150
FedFS has been deprecated	150
Btrfs has been deprecated	150
tcp_wrappers deprecated	151
nautilus-open-terminal replaced with gnome-terminal-nautilus	151
sslwrap() removed from Python	151
Symbols from libraries linked as dependencies no longer resolved by ld	151
Windows guest virtual machine support limited	151
libnetlink is deprecated	151
S3 and S4 power management states for KVM have been deprecated	151
The Certificate Server plug-in udnPwdDirAuth is discontinued	151
Red Hat Access plug-in for IdM is discontinued	152
The Ipsilon identity provider service for federated single sign-on	152
Several rsyslog options deprecated	152
Deprecated symbols from the memkind library	152
Options of Sockets API Extensions for SCTP (RFC 6458) deprecated	153
Managing NetApp ONTAP using SSLv2 and SSLv3 is no longer supported by libstorageMgmt	153
dconf-dbus-1 has been deprecated and dconf-editor is now delivered separately	153
FreeRADIUS no longer accepts Auth-Type := System	153
Deprecated Device Drivers	153
SFN4XXX adapters have been deprecated	156
Software initiated only FCoE storage technologies have been deprecated	156
Containers using the libvirt-lxc tooling have been deprecated	156
PART V. KNOWN ISSUES	157
CHAPTER 54. AUTHENTICATION AND INTEROPERABILITY	158
sudo unexpectedly denies access when performing group lookups	158
The KCM credential cache is not suitable for a large number of credentials in a single credential cache	158
The sssd-secrets component crashes when it is under load	158
SSSD does not correctly handle multiple certificate matching rules with the same priority	159
SSSD can look up only unique certificates in ID overrides	159

The ipa-adviser command does not fully configure smart card authentication	159
The libwbclient library fails to connect to Samba shares hosted on Red Hat Enterprise Linux 7.4	159
Certificate System subsystems experience communication problems with TLS_ECDHE_RSA_* ciphers and certain HSMTs	159
CHAPTER 55. COMPILER AND TOOLS	160
Performance of regular expressions cannot be boosted with the JIT technique if executable stack is disabled	160
Memory leaks occur when certain applications fail to exit after unloading the Gluster libraries	160
URL to DISA SRGs is incorrect	160
The ensure_gpgcheck_repo_metadata rule fails	160
The SSG pam_faillock module utilization check incorrectly accepts default=die	160
CHAPTER 56. DESKTOP	161
Updating totem alone fails	161
The operating system always assumes Wacom Expresskeys Remote mode 1 when booting	161
Cannot install downloaded RPM files from Nautilus	161
Yelp does not correctly display HTML formatted files	161
Automatic modesetting fails when attaching monitors with some AMD hardware	161
Gnome Documents can not display some documents when installed without LibreOffice due to a missing dependency	161
Application Installer displays packages even though they can not be installed on big endian architectures	162
The Add/Remove Software tool (gpk-application) does not use a newly imported key on the first try	162
Resizing a display of a virtual machine with multiple displays using multiple PCI devices causes X to crash	162
Nautilus does not hide icons in the GNOME Classic Session	162
Incorrect dependency in flatpak	162
Firefox does not start after update	162
Limited support for visuals in Xorg	163
CHAPTER 57. FILE SYSTEMS	164
NetApp storage appliances serving NFSv4 are advised to check their configuration	164
CHAPTER 58. HARDWARE ENABLEMENT	165
The i40e driver rejects the most general HWTSTAMP filter	165
CHAPTER 59. INSTALLATION AND BOOTING	166
FIPS mode unsupported when installing from an HTTPS kickstart source	166
PXE boot with UEFI and IPv6 displays the GRUB2 shell instead of the operating system selection menu	166
Specifying a driverdisk partition with non-alphanumeric characters generates an invalid output Kickstart file	166
The Scientific Computing variant is missing packages required for certain security profiles	166
CHAPTER 60. KERNEL	167
kexec fails when secondary cores do not offline	167
File-system corruption due to incorrect flushing of cache has been fixed but I/O operations can be slower	167
Wacom Cintiq 12WX is not redetected when unplugged and plugged in quickly	167
Installing to some IBM POWER8 machines using a Virtual DVD fails when starting GUI	167
Entering full screen mode using a keyboard shortcut causes display problems on VMWare ESXi 5.5	167
KSC currently does not support xz compression	168
CHAPTER 61. NETWORKING	169
Verification of signatures using the MD5 hash algorithm is disabled in Red Hat Enterprise Linux 7	169
CHAPTER 62. SECURITY	170
certutil does not return the NSS database password requirements in FIPS mode	170
systemd-importd runs as init_t	170

The SCAP password length requirement is ignored in the kickstart installation	170
rhnsd.pid is writable by group and others	170
CHAPTER 63. STORAGE	171
No support for thin provisioning on top of RAID in a cluster	171
Anaconda installation can fail when LVM or md device has metadata from a previous install	171
CHAPTER 64. SYSTEM AND SUBSCRIPTION MANAGEMENT	172
System upgrade may cause Yum to install unneeded 32-bit packages if rdma-core is installed	172
CHAPTER 65. VIRTUALIZATION	173
Booting OVMF guests fails	173
Bridge creation with virsh iface-bridge fails	173
Guests sometimes fail to boot on ESXi 5.5	173
The STIG for Red Hat Virtualization Hypervisor profile is not displayed in Anaconda	173
APPENDIX A. COMPONENT VERSIONS	174
APPENDIX B. LIST OF BUGZILLAS BY COMPONENT	175
APPENDIX C. REVISION HISTORY	189

PREFACE

Red Hat Enterprise Linux minor releases are an aggregation of individual security, enhancement, and bug fix errata. The *Red Hat Enterprise Linux 7.4 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

- Red Hat Enterprise Linux 7.4 introduces support for Network Bound Disk Encryption (NBDE), which enables the system administrator to encrypt root volumes of hard drives on bare metal machines without requiring to manually enter password when systems are rebooted.
- The **USBGuard** software framework provides system protection against intrusive USB devices by implementing basic whitelisting and blacklisting capabilities based on device attributes.
- The **OpenSSH** libraries update includes the ability to resume interrupted uploads in Secure File Transfer Protocol (SFTP) and adds support for a new fingerprint type that uses the SHA-256 algorithm. This **OpenSSH** version also removes server-side support for the SSH-1 protocol.
- Multiple new Linux Audit capabilities have been added to enable easier administration, to filter the events logged by the Audit system, gather more information from critical events, and to interpret large numbers of records.
- The **OpenSC** set of libraries and utilities adds support for Common Access Card (CAC) cards and now provides also the **CoolKey** applet functionality.
- The **OpenSSL** update includes multiple enhancements, such as support for the Datagram Transport Layer Security (DTLS) version 1.2 protocol and Application-Layer Protocol Negotiation (ALPN).
- The **OpenSCAP** tools have been NIST-certified, which enables easier adoption in regulated environments.
- Cryptographic protocols and algorithms that are considered insecure have been deprecated. However, this version also introduces a lot of other cryptographic-related improvements. For more information, see [Chapter 53, *Deprecated Functionality*](#) and the [Enhancing the Security of the Operating System with Cryptography Changes in Red Hat Enterprise Linux 7.4](#) Knowledgebase article on the Red Hat Customer Portal.

See [Chapter 15, *Security*](#) for more information on security enhancements.

Identity Management

- The System Security Services Daemon (SSSD) in a container is now fully supported. The Identity Management (IdM) server container is available as a Technology Preview feature.
- Users are now able to install new Identity Management servers, replicas, and clients on systems with FIPS mode enabled.
- Several enhancements related to smart card authentication have been introduced.

For detailed information on changes in IdM, see [Chapter 5, *Authentication and Interoperability*](#). For details on deprecated capabilities related to IdM, see [Chapter 53, *Deprecated Functionality*](#).

Networking

- **NetworkManager** supports additional features for routing, enables the Media Access Control Security (MACsec) technology, and is now able to handle unmanaged devices.
- Kernel Generic Routing Encapsulation (GRE) tunneling has been enhanced.

For more networking features, see [Chapter 14, Networking](#).

Kernel

- Support for NVMe Over Fabric has been added to the NVM-Express kernel driver, which increases flexibility when accessing high performance NVMe storage devices located in the data center on both Ethernet or Infiniband fabric infrastructures.

For further kernel-related changes, refer to [Chapter 12, Kernel](#).

Storage and File Systems

- LVM provides full support for RAID takeover, which allows users to convert a RAID logical volume from one RAID level to another, and for RAID reshaping, which allows users to reshape properties, such as the RAID algorithm, stripe size, or number of images.
- You can now enable SELinux support for containers when you use OverlayFS with Docker.
- NFS over RDMA (NFSv4.1) server is now fully supported when accessed by Red Hat Enterprise Linux clients.

See [Chapter 17, Storage](#) for further storage-related features and [Chapter 9, File Systems](#) for enhancements to file systems.

Tools

- The **Performance Co-Pilot (PCP)** application has been enhanced to support new client tools, such as `pcp2influxdb`, `pcp-mpstat`, and `pcp-pidstat`. Additionally, new PCP performance metrics from several subsystems are available for a variety of Performance Co-Pilot analysis tools.

For more information regarding updates to various tools, see [Chapter 7, Compiler and Tools](#).

High Availability

- Red Hat Enterprise Linux 7.4 introduces full support for the following features:
 - `cluftr`, a tool for transforming and analyzing cluster configuration formats
 - Quorum devices (QDevice) in a Pacemaker cluster for managing stretch clusters
 - `Booth` cluster ticket manager

For more information on the high availability features introduced in this release, see [Chapter 6, Clustering](#).

Virtualization

- Red Hat Enterprise Linux 7 guest virtual machines now support the Elastic Network Adapter (ENA), and thus provide enhanced networking capabilities when running on the the Amazon Web Services (AWS) cloud.

For further enhancements to Virtualization, see [Chapter 19, Virtualization](#).

Management and Automation

- Red Hat Enterprise Linux 7.4 includes **Red Hat Enterprise Linux System Roles** powered by **Ansible**, a configuration interface that simplifies management and maintenance of Red Hat Enterprise Linux deployments. This feature is available as a Technology Preview.

Enterprise Linux deployments. This feature is available as a Technology Preview.

For details, refer to [Chapter 47, Red Hat Enterprise Linux System Roles Powered by Ansible](#)

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

Red Hat Customer Portal Labs

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

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

CHAPTER 2. ARCHITECTURES

Red Hat Enterprise Linux 7.4 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]

[1] Note that the Red Hat Enterprise Linux 7.4 installation is supported only on 64-bit hardware. Red Hat Enterprise Linux 7.4 is able to run 32-bit operating systems, including previous versions of Red Hat Enterprise Linux, as virtual machines.

[2] Red Hat Enterprise Linux 7.4 (big endian) is currently supported as a KVM guest on Red Hat Enterprise Virtualization for Power, and on PowerVM.

[3] Red Hat Enterprise Linux 7.4 (little endian) is currently supported as a KVM guest on Red Hat Enterprise Virtualization for Power, on PowerVM, and PowerNV (bare metal).

[4] Note that Red Hat Enterprise Linux 7.4 supports IBM zEnterprise 196 hardware or later; IBM z10 Systems mainframe systems are no longer supported and will not boot Red Hat Enterprise Linux 7.4.

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 7.4. These changes include added or updated `proc` entries, `sysctl`, and `sysfs` default values, boot parameters, kernel configuration options, or any noticeable behavior changes.

UPDATED /PROC/SYS/KERNEL ENTRIES

`hung_task_panic`

Controls the behavior of the kernel when an unresponsive task is detected. This file occurs if `CONFIG_DETECT_HUNG_TASK` is enabled.

Format: { "0" | "1" }

0 - Continue operation. Default behavior.

1 - Panic immediately.

`hung_task_check_count`

Provides the upper bound on the number of tasks that are checked. This file occurs if `CONFIG_DETECT_HUNG_TASK` is enabled.

`hung_task_timeout_secs`

Checks interval. Reports a warning in case that a task in D state is not scheduled for longer time than this value. This file occurs if `CONFIG_DETECT_HUNG_TASK` is enabled.

0 - Infinite timeout - no checking done.

`hung_task_warning`

Provides the maximum number of warnings to report during a check interval. When this value is reached, no more warnings will be reported. This file occurs if `CONFIG_DETECT_HUNG_TASK` is enabled.

-1 - Reports an infinite number of warnings.

`panic_on_rcu_stall`

When set to 1, calls the `panic()` function after RCU stall detection messages. This is useful to define the root cause of RCU stalls using a vmcore.

0 - Do not panic when RCU stall takes place. Default behavior.

1 - Panic after printing RCU stall messages.

UPDATED /PROC/SYS/USER ENTRIES

You can use the files in the `/proc/sys/user` directory to override the default limits for the number of namespaces and other objects that have per user namespace limits. The purpose of these limits is to stop programs that malfunction and attempt to create a high number of objects. The default values of these limits are adjusted so that any program in normal operation cannot reach them.

The creation of per user namespace objects is charged to the user in the user namespace who created the object and who verified to be below the per user limit in that user namespace. The creation of such objects happens in user namespaces and is also charged to all users who created user namespaces.

This recursive counting of created objects ensures that creating a user namespace does not allow a user to exceed their current limits.

The updated files in `/proc/sys/user` are:

max_cgroup_namespaces

The maximum number of control group namespaces that any user in the current user namespace can create.

max_ipc_namespaces

The maximum number of interprocess communication namespaces that any user in the current user namespace can create.

max_mnt_namespaces

The maximum number of mount namespaces that any user in the current user namespace can create.

max_net_namespaces

The maximum number of network namespaces that any user in the current user namespace can create.

max_pid_namespaces

The maximum number of process ID namespaces that any user in the current user namespace can create.

max_user_namespaces

The maximum number of user ID namespaces that any user in the current user namespace can create.

max_uts_namespaces

The maximum number of UNIX Timesharing System (UTS) namespaces that any user in the current user namespace can create.

KERNEL PARAMETERS

acpi_force_table_verification [HW,ACPI]

Enables table checksum verification during early stage. By default, disabled on the 32-bit AMD and Intel architecture due to early mapping size limitation.

acpi_no_auto_ssdt [HW,ACPI]

Disables automatic loading of Secondary System Description Table (SSDT).

acpi_no_static_ssdt [HW,ACPI]

Disables installation of static SSDTs at early boot time. By default, SSDTs contained in the Root System Description Table (RSDT) or eXtended System Descriptor Table (XSDT) are installed automatically and they appear in the `/sys/firmware/acpi/tables` directory.

This option turns off this feature. Specifying this option does not affect dynamic table installation which installs SSDT tables to the `/sys/firmware/acpi/tables/dynamic` directory.

irqaffinity= [SMP]

Sets the default irq affinity mask in the following formats:

Format: `<cpu number>,..., <cpu number>`

or

`<cpu number>-<cpu number>`

You can use a positive range in ascending order or a combination.

`<cpu number>,...,<cpu number>-<cpu number>`

nokaslr [KNL]]

Disables installation of static SSDTs at early boot time. By default, SSDTs contained in the RSDT or XSDT are installed automatically and they appear in the `/sys/firmware/acpi/tables` directory.

Disables kernel and module base offset Address SpaceLayout Randomization (ASLR) if `CONFIG_RANDOMIZE_BASE` is set.

nohibernate

Disables hibernation and resume.

crash_kexec_post_notifiers

Runs `kdump` after running panic-notifiers and dumping kmsg.

[PCI] hpbussize=nn

Provides the minimum amount of additional bus numbers reserved for buses below a hotplug bridge. Default is 1.

pcie_port_pm=[PCIE]

PCIe port power management handling:

Format: { "off" | "force" }

off - Disables power management of all PCIe ports.

1 - Enables power management of all PCIe ports.

sunrpc.svc_rpc_per_connection_limit=[NFS,SUNRPC]

Limits the number of requests for the server to process in parallel from a single connection. The default value is 0 (no limit).

PART I. NEW FEATURES

This part documents new features and major enhancements introduced in Red Hat Enterprise Linux 7.4.

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)

cloud-init moved to the Base channel

As of Red Hat Enterprise Linux 7.4, the cloud-init package and its dependencies have been moved from the Red Hat Common channel to the Base channel. **Cloud-init** is a tool that handles early initialization of a system using metadata provided by the environment. It is typically used to configure servers booting in a cloud environment, such as OpenStack or Amazon Web Services. Note that the cloud-init package has not been updated since the latest version provided through the Red Hat Common channel. (BZ#1427280)

CHAPTER 5. AUTHENTICATION AND INTEROPERABILITY

SSSD in a container now fully supported

The `rhel7/sssd` container image, which provides the System Security Services Daemon (SSSD), is no longer a Technology Preview feature. The image is now fully supported. Note that the `rhel7/ipa-server` container image is still a Technology Preview feature.

For details, see https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html-single/using_containerized_identity_management_services. (BZ#1467260)

Identity Management now supports FIPS

With this enhancement, Identity Management (IdM) supports the Federal Information Processing Standard (FIPS). This enables you to run IdM in environments that must meet the FIPS criteria. To run IdM with FIPS mode enabled, you must set up all servers in the IdM environment using Red Hat Enterprise Linux 7.4 with FIPS mode enabled.

Note that you cannot:

- Enable FIPS mode on existing IdM servers previously installed with FIPS mode disabled.
- Install a replica in FIPS mode when using an existing IdM server with FIPS mode disabled.

For further details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html-single/Linux_Domain_Identity_Authentication_and_Policy_Guide/index.html#prerequisites. (BZ#1125174)

SSSD supports obtaining a Kerberos ticket when users authenticate with a smart card

The System Security Services Daemon (SSSD) now supports the Kerberos PKINIT preauthentication mechanism. When authenticating with a smart card to a desktop client system enrolled in an Identity Management (IdM) domain, users receive a valid Kerberos ticket-granting ticket (TGT) if the authentication was successful. Users can then use the TGT for further single sign-on (SSO) authentication from the client system.

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/sc-pkinit-auth.html. (BZ#1200767, BZ#1405075)

SSSD enables logging in to different user accounts with the same smart card certificate

Previously, the System Security Services Daemon (SSSD) required every certificate to be uniquely mapped to a single user. When using smart card authentication, users with multiple accounts were not able to log in to all of these accounts with the same smart card certificate. For example, a user with a personal account and a functional account (such as a database administrator account) was able to log in only to the personal account.

With this update, SSSD no longer requires certificates to be uniquely mapped to a single user. As a result, users can now log in to different accounts with a single smart card certificate.

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/smart_cards.html. (BZ#1340711, BZ#1402959)

IdM web UI enables smart card login

The Identity Management web UI enables users to log in using smart cards.

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/sc-web-ui-auth.html. (BZ#1366572)

New packages: keycloak-httpd-client-install

The keycloak-httpd-client-install packages provide various libraries and tools that can automate and simplify the configuration of Apache `httpd` authentication modules when registering as a Red Hat Single Sign-On (RH-SSO, also called Keycloak) federated Identity Provider (IdP) client.

For details on RH-SSO, see <https://access.redhat.com/products/red-hat-single-sign-on>.

As part of this update, new dependencies have been added to Red Hat Enterprise Linux:

- The `python-requests-oauthlib` package: This package provides the OAuth library support for the `python-requests` package, which enables `python-requests` to use OAuth for authentication.
- The `python-oauthlib` package: This package is a Python library providing OAuth authentication message creation and consumption. It is meant to be used in conjunction with tools providing message transport. (BZ#1401781, BZ#1401783, BZ#1401784)

New Kerberos credential cache type: KCM

This update adds a new SSSD service named `kcm`. The service is included in the `sssd-kcm` subpackage.

When the `kcm` service is installed, you can configure the Kerberos library to use a new credential cache type named **KCM**. When the KCM credential cache type is configured, the `sssd-kcm` service manages the credentials.

The KCM credential cache type is well-suited for containerized environments:

- With KCM, you can share credential caches between containers on demand, based on mounting the UNIX socket on which the `kcm` service listens.
- The `kcm` service runs in user space outside the kernel, unlike the **KEYRING** credential cache type that RHEL uses by default. With KCM, you can run the `kcm` service only in selected containers. With **KEYRING**, all containers share the credential caches because they share the kernel.

Additionally, the KCM credential cache type supports cache collections, unlike the **FILE** `ccache` type.

For details, see the `sssd-kcm(8)` man page. (BZ#1396012)

AD users can log in to the web UI to access their self-service page

Previously, Active Directory (AD) users were only able to authenticate using the `kindit` utility from the command line. With this update, AD users can also log in to the Identity Management (IdM) web UI. Note that the IdM administrator must create an ID override for an AD user before the user is able to log in.

As a result, AD users can access their self-service page through the IdM web UI. The self-service page displays the information from the AD users' ID override.

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/using-the-ui.html#ad-users-idm-web-ui. (BZ#872671)

SSSD enables configuring an AD subdomain in the SSSD server mode

Previously, the System Security Services Daemon (SSSD) automatically configured trusted Active Directory (AD) domains. With this update, SSSD supports configuring certain parameters for trusted AD domains in the same way as the joined domain.

As a result, you can set individual settings for trusted domains, such as the domain controller that SSSD communicates with. To do this, create a section in the `/etc/sss/sss.conf` file with a name that follows this template:

```
[domain/main_domain/trusted_domain]
```

For example, if the main IdM domain name is `ipa.com` and the trusted AD domain name is `ad.com`, the corresponding section name is:

```
[domain/ipa.com/ad.com]
```

(BZ#1214491)

SSSD supports user and group lookups and authentication with short names in AD environments

Previously, the System Security Services Daemon (SSSD) supported user names without the domain component, also called short names, for user and group resolution and authentication only when the daemon was joined to a standalone domain. Now, you can use short names for these purposes in all SSSD domains in these environments:

- On clients joined to Active Directory (AD)
- In Identity Management (IdM) deployments with a trust relationship to an AD forest

The output format of all commands is always fully-qualified even when using short names. This feature is enabled by default after you set up a domain's resolution order list in one of the following ways (listed in order of preference):

- Locally, by configuring the list using the `domain_resolution_order` option in the `[sss]` section of the `/etc/sss/sss.conf` file
- By using an ID view
- Globally, in the IdM configuration

To disable the feature, set the `use_fully_qualified_names` option to `True` in the `[domain/example.com]` section of the `/etc/sss/sss.conf` file. (BZ#1330196)

SSSD supports user and group resolution, authentication, and authorization in setups without UIDs or SIDs

In traditional System Security Services Daemon (SSSD) deployments, users and groups either have POSIX attributes set or SSSD can resolve the users and groups based on Windows security identifiers (SID).

With this update, in setups that use LDAP as the identity provider, SSSD now supports the following functionality even when UIDs or SIDs are not present in the LDAP directory:

- User and group resolution through the D-Bus interface
- Authentication and authorization through the pluggable authentication module (PAM) interface (BZ#1425891)

SSSD introduces the `sssctl user-checks` command, which checks basic SSSD functionality in a single operation

The `sssctl` utility now includes a new command named `user-checks`. The `sssctl user-checks` command helps debug problems in applications that use the System Security Services Daemon (SSSD) as a back end for user lookup, authentication, and authorization.

- The `sssctl user-checks [USER_NAME]` command displays user data available through Name Service Switch (NSS) and the InfoPipe responder for the D-Bus interface. The displayed data shows whether the user is authorized to log in using the `system-auth` pluggable authentication module (PAM) service.
- Additional options accepted by `sssctl user-checks` check authentication or different PAM services.

For details on `sssctl user-checks`, use the `sssctl user-checks --help` command. (BZ#1414023)

Support for secrets as a service

This update adds a responder named `secrets` to the System Security Services Daemon (SSSD). This responder allows an application to communicate with SSSD over a UNIX socket using the Custodia API. This enables SSSD to store secrets in its local database or to forward them to a remote Custodia server. (BZ#1311056)

IdM enables semi-automatic upgrades of the IdM DNS records on an external DNS server

To simplify updating the Identity Management (IdM) DNS records on an external DNS server, IdM introduces the `ipa dns-update-system-records --dry-run --out [file]` command. The command generates a list of records in a format accepted by the `nsupdate` utility.

You can use the generated file to update the records on the external DNS server by using a standard dynamic DNS update mechanism secured with the Transaction Signature (TSIG) protocol or the GSS algorithm for TSIG (GSS-TSIG).

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/dns-updates-external.html. (BZ#1409628)

IdM now generates SHA-256 certificate and public key fingerprints

Previously, Identity Management (IdM) used the MD5 hash algorithm when generating fingerprints for certificates and public keys. To increase security, IdM now uses the SHA-256 algorithm in the mentioned scenario. (BZ#1444937)

IdM supports flexible mapping mechanisms for linking smart card certificates to user accounts

Previously, the only way to find a user account corresponding to a certain smart card in Identity Management (IdM) was to provide the whole smart card certificate as a Base64-encoded DER string. With this update, it is possible to find a user account also by specifying attributes of the smart card certificates, not just the certificate string itself. For example, the administrator can now define matching and mapping rules to link smart card certificates issued by a certain certificate authority (CA) to a user account in IdM.

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/smart_cards.html#sc-one-card-multiple-accounts-links. (BZ#1402959)

New user-space tools enable a more convenient LMDB debugging

This update introduces the `mdb_copy`, `mdb_dump`, `mdb_load`, and `mdb_stat` tool in the `/usr/libexec/openldap/` directory. The addition includes relevant man pages in the `man/man1` subdirectory. Use the new tools only to debug problems related to the Lightning Memory-Mapped Database (LMDB) back end. (BZ#1428740)

openldap rebased to version 2.4.44

The openldap packages have been upgraded to upstream version 2.4.44, which provides a number of bug fixes and enhancements over the previous version. In particular, this new version fixes many replication and Lightning Memory-Mapped Database (LMDB) bugs. (BZ#1386365)

Improved security of DNS lookups and robustness of service principal lookups in Identity Management

The Kerberos client library no longer attempts to canonicalize host names when issuing ticket-granting server (TGS) requests. This feature improves:

- Security because DNS lookups, which were previously required during canonicalization, are no longer performed
- Robustness of service principal lookups in more complex DNS environments, such as clouds or containerized applications

Make sure you specify the correct fully qualified domain name (FQDN) in host and service principals. Due to this change in behavior, Kerberos does not attempt to resolve any other form of names in principals, such as short names. (BZ#[1404750](#))

samba rebased to version 4.6.2

The samba packages have been upgraded to version 4.6.2, which provides a number of bug fixes and enhancements over the previous version:

- Samba now verifies the ID mapping configuration before the `winbindd` service starts. If the configuration is invalid, `winbindd` fails to start. Use the `testparm` utility to verify your `/etc/samba/smb.conf` file. For further details, see the **IDENTITY MAPPING CONSIDERATIONS** section in the `smb.conf` man page.
- Uploading printer drivers from Windows 10 now works correctly.
- 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.
- The `net ads unregister` command can now delete the DNS entry of the host from the Active Directory DNS zone when leaving the domain.
- SMB 2.1 leases are now enabled by default in the `smb2 leases` parameter. SMB leasing enables clients to aggressively cache files.
- To improve security, the NT LAN manager version 1 (NTLMv1) protocol is now disabled by default. If you require the insecure NTLMv1 protocol, set the `ntlm auth` parameter in the `/etc/samba/smb.conf` file to `yes`.
- The `event` subcommand has been added to the `ctdb` utility for interacting with event scripts.
- The `idmap_hash` ID mapping back end is marked as deprecated will be removed in a future Samba version.

- The deprecated `only` `user` and `username` parameters have been removed.

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

For further information about notable changes, read the upstream release notes before updating. (BZ#1391954)

authconfig can enable SSSD to authenticate users with smart cards

This new feature allows the `authconfig` command to configure the System Security Services Daemon (SSSD) to authenticate users with smart cards, for example:

```
# authconfig --enablesssd --enablesssdauth --enablesmartcard --
smartcardmodule=sssd --smartcardaction=0 --updateall
```

With this update, smart card authentication can now be performed on systems where `pam_pkcs11` is not installed. However, if `pam_pkcs11` is installed, the `--smartcardmodule=sssd` option is ignored. Instead, the first `pkcs11_module` defined in the `/etc/pam_pkcs11/pam_pkcs11.conf` is used as default.

For details, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/auth-idm-client-sc.html. (BZ#1378943)

authconfig can now enable account locking

This update adds the `--enablefaillock` option for the `authconfig` command. When the option is enabled, the configured account will be locked for 20 minutes after four consecutive failed login attempts within a 15-minute interval. (BZ#1334449)

Improved performance of the IdM server

The Identity Management (IdM) server has a higher performance across many of the common workflows and setups. These improvements include:

- Vault performance has been increased by reducing the round trips within the IdM server management framework.
- The IdM server management framework has been tuned to reduce the time spent in internal communication and authentication.
- The Directory Server connection management has been made more scalable with the use of the `nunc-stans` framework.
- On new installations, the Directory Server now auto-tunes the database entry cache and the number of threads based on the hardware resources of the server.
- The `memberOf` plug-in performance has been improved when working with large or nested groups. (BZ#1395940, BZ#1425906, BZ#1400653)

The default session expiration period in the IdM web UI has changed

Previously, when the user logged in to the Identity Management (IdM) web UI using a user name and password, the web UI automatically logged the user out after 20 minutes of inactivity. With this update, the default session length is the same as the expiration period of the Kerberos ticket obtained during the login operation. To change the default session length, use the `kinit_lifetime` option in the `/etc/ipa/default.conf` file, and restart the `httpd` service. (BZ#1459153)

The `dbmon.sh` script now uses instance names to connect to Directory Server instances

The `dbmon.sh` shell script enables you to monitor the Directory Server database and entry cache usage. With this update, the script no longer uses the `HOST` and `PORT` environment variables. To support secure binds, the script now reads the Directory Server instance name from the `SERVID` environment variable and uses it to retrieve the host name, port, and the information if the server requires a secure connection. For example, to monitor the `slapd-localhost` instance, enter:

```
SERVID=slapd-localhost INCR=1 BINDDN="cn=Directory Manager"  
BINDPW="password" dbmon.sh
```

(BZ#[1394000](#))

Directory Server now uses the SSHA_512 password storage scheme as default

Previously, Directory Server used the weak 160-bit salted secure hash algorithm (SSHA) as default password storage scheme set in the `passwordStorageScheme` and `nsslapd-rootpwstoragescheme` parameters in the `cn=config` entry. To increase security, the default of both parameters has been changed to the strong 512-bit SSHA scheme (SSHA_512).

The new default is used:

- When performing new Directory Server installations.
- When the `passwordStorageScheme` parameter is not set, and you are updating passwords stored in `userPassword` attributes.
- When the `nsslapd-rootpwstoragescheme` parameter is not set, and you are updating the Directory Server manager password set in the `nsslapd-rootpw` attribute. (BZ#[1425907](#))

Directory Server now uses the `tcmalloc` memory allocator

Red Hat Directory Server now uses the `tcmalloc` memory allocator. The previously used standard `glibc` allocator required more memory, and in certain situations, the server could run out of memory. Using the `tcmalloc` memory allocator, Directory Server now requires less memory, and the performance increased. (BZ#[1426275](#))

Directory Server now uses the `nunc-stans` framework

The `nunc-stans` event-based framework has been integrated into Directory Server. Previously, the performance could be slow when many simultaneous incoming connections were established to Directory Server. With this update, the server is able to handle a significantly larger number of connections without performance degradation. (BZ#[1426278](#), BZ#[1206301](#), BZ#[1425906](#))

Improved performance of the Directory Server `memberof` plug-in

Previously, when working with large or nested groups, plug-in operations could take a long time. With this update, the performance of the Red Hat Directory Server `memberof` plug-in has been improved. As a result, the `memberof` plug-in now adds and removes users faster from groups. (BZ# [1426283](#))

Directory Server now logs severity levels in the error log file

Directory Server now logs severity levels in the `/var/log/dirsrv/slapd-instance_name/errors` log file. Previously, it was difficult to distinguish the severity of entries in the error log file. With this enhancement, administrators can use the severity level to filter the error log.

For further details, see the corresponding section in the Red Hat Directory Server Configuration, Command, and File Reference: <https://access.redhat.com/documentation/en->

[US/Red_Hat_Directory_Server/10/html/Configuration_Command_and_File_Reference/error-logs.html#error-logs-content](https://access.redhat.com/documentation/en-US/Red_Hat_Directory_Server/10/html/Configuration_Command_and_File_Reference/error-logs.html#error-logs-content) (BZ#1426289)

Directory Server now supports the PBKDF2_SHA256 password storage scheme

To increase security, this update adds the 256-bit password-based key derivation function 2 (PBKDF2_SHA256) to the list of supported password-storage schemes in Directory Server. The scheme uses 30,000 iterations to apply the 256-bit secure hash algorithm (SHA256).

Note that the network security service (NSS) database in Red Hat Enterprise Linux prior to version 7.4 does not support PBKDF2. Therefore, you cannot use this password scheme in a replication topology with previous Directory Server versions. (BZ#1436973)

Improved auto-tuning support in Directory Server

Previously, you had to monitor the databases and manually tune settings to improve the performance. With this update, Directory Server supports optimized auto-tuning for:

- The database and entry cache
- The number of threads created

Directory Server tunes these settings, based on the hardware resources of the server.

Auto-tuning is now automatically enabled by default if you install a new Directory Server instance. On instances upgraded from earlier versions, Red Hat recommends to enable auto-tuning. For details, see:

- Database and entry cache: https://access.redhat.com/documentation/en-US/Red_Hat_Directory_Server/10/html/Performance_Tuning_Guide/memoryusage.html#DB_ar
- Directory Server threads: https://access.redhat.com/documentation/en-US/Red_Hat_Directory_Server/10/html/Performance_Tuning_Guide/ds-threads (BZ#1426286)

New PKI configuration parameter allows control of the TCP keepAlive option

This update adds the `tcp.keepAlive` parameter to the `CS.cfg` configuration file. This parameter accepts boolean values, and is set to `true` by default. Use this parameter to configure the TCP `keepAlive` option for all LDAP connections created by the PKI subsystem. This option is useful in cases where certificate issuance takes a very long time and connections are being closed automatically after being idle for too long. (BZ#1413132)

PKI Server now creates PKCS #12 files using strong encryption

When generating PKCS #12 files, the `pki pkcs12` command previously used the PKCS #12 deprecated key derivation function (KDF) and the triple DES (3DES) algorithm. With this update, the command now uses the password-based encryption standard 2 (PBES2) scheme with the password-based key derivation function 2 (PBKDF2) and the Advanced Encryption Standard (AES) algorithm to encrypt private keys. As a result, this enhancement increases the security and complies the Common Criteria certification requirements. (BZ#1426754)

CC-compliant algorithms available for encryption operations

Common Criteria requires that encryption and key-wrapping operations are performed using approved algorithms. These algorithms are specified in section FCS_COP.1.1(1) in the Protection Profile for Certification Authorities. This update modifies encryption and decryption in the KRA to use approved AES encryption and wrapping algorithms in the transport and storage of secrets and keys. This update required changes in both the server and client software. (BZ#1445535)

New options to allow configuring visibility of menu items in the TPS interface

Previously, menu items grouped under the **System** menu in the Token Processing System (TPS) user

interface were determined statically based on user roles. In certain circumstances, the displayed menu items did not match components actually accessible by the user. With this update, the `System` menu in the TPS user interface only displays menu items based on the `target.configure.list` parameter for TPS administrators, and the `target.agent_approve.list` parameter for TPS agents. These parameters can be modified in the instance `CS.cfg` file to match accessible components. (BZ#1391737)

Added a profile component to copy certificate Subject Common Name to the Subject Alternative Name extension

Some TLS libraries now warn or refuse to validate DNS names when the DNS name only appears in the Subject Common Name (CN) field, which is a practice that was deprecated by RFC 2818. This update adds the `CommonNameToSANDefault` profile component, which copies the Subject Common Name to the Subject Alternative Name (SAN) extension, and ensures that certificates are compliant with current standards. (BZ#1305993)

New option to remove LDAP entries before LDIF import

When migrating a CA, if an LDAP entry existed before the LDIF import, then the entry was not recreated from the LDAP import, causing some fields to be missing. Consequently, the request ID showed up as undefined. This update adds an option to remove the LDAP entry for the signing certificate at the end of the `pkispawn` process. This entry is then re-created in the subsequent LDIF import. Now, the request ID and other fields show up correctly if the signing entry is removed and re-added in the LDIF import. The correct parameters to add are (X represents the serial number of the signing certificate being imported, in decimal):

```
pki_ca_signing_record_create=False
pki_ca_signing_serial_number=X
```

(BZ#1409946)

Certificate System now supports externally authenticated users

Previously, you had to create users and roles in Certificate System. With this enhancement, you can now configure Certificate System to admit users authenticated by an external identity provider. Additionally, you can use realm-specific authorization access control lists (ACLs). As a result, it is no longer necessary to create users in Certificate System. (BZ#1303683)

Certificate System now supports enabling and disabling certificate and CRL publishing

Prior to this update, if publishing was enabled in a certificate authority (CA), Certificate System automatically enabled both certificate revocation list (CRL) and certificate publishing. Consequently, on servers that did not have certificate publishing enabled, error messages were logged. Certificate System has been enhanced, and now supports enabling and disabling certificate and CRL publishing independently in the `/var/lib/pki/<instance>/ca/conf/CS.cfg` file.

To enable or disable both certificate and CRL publishing, set:

```
ca.publish.enable = True|False
```

To enable only CRL publishing, set:

```
ca.publish.enable = True
ca.publish.cert.enable = False
```

To enable only certificate publishing, set:

```
ca.publish.enable = True
ca.publish.crl.enable = False
```

(BZ#[1325071](#))

The searchBase configuration option has been added to the DirAc1Authz PKI Server plug-in

To support reading different sets of authorization access control lists (ACL), the `searchBase` configuration option has been added to the `DirAc1Authz` PKI Server plug-in. As a result, you can set the sub-tree from which the plug-in loads ACLs. (BZ#[1388622](#))

For better performance, Certificate System now supports ephemeral

Before this update, Certificate System key recovery agent (KRA) instances always stored recovery and storage requests of secrets in the LDAP back end. This is required to store the state if multiple agents must approve the request. However, if the request is processed immediately and only one agent must approve the request, storing the state is not required. To improve performance, you can now set the `kra.ephemeralRequests=true` option in the `/var/lib/pki/<instance>/kra/conf/CS.cfg` file to no longer store requests in the LDAP back end. (BZ#[1392068](#))

Section headers in PKI deployment configuration file are no longer case sensitive

The section headers (such as `[Tomcat]`) in the PKI deployment configuration file were previously case-sensitive. This behavior increased the chance of an error while providing no benefit. Starting with this release, section headers in the configuration file are case-insensitive, reducing the chance of an error occurring. (BZ#[1447144](#))

Certificate System now supports installing a CA using HSM on FIPS-enabled Red Hat Enterprise Linux

During the installation of a Certificate System Certificate Authority (CA) instance, the installer needs to restart the instance. During this restart, instances on an operating system having the Federal Information Processing Standard (FIPS) mode enabled and using a hardware security module (HSM), need to connect to the non-secure HTTP port instead of the HTTPS port. With this update, it is now possible to install a Certificate System instance on FIPS-enabled Red Hat Enterprise Linux using an HSM. (BZ#[1450143](#))

CMC requests now use a random IV for AES and 3DES encryption

With this update, Certificate Management over CMS (CMC) requests in PKI Server use a randomly generated initialization vector (IV) when encrypting a key to be archived. Previously, the client and server code used a fixed IV in this scenario. The CMC client code has been enhanced, and as a result, using random IVs increase security when performing encryption for both Advanced Encryption Standard (AES) and Triple Data Encryption Algorithm (3DES). (BZ#[1458055](#))

CHAPTER 6. CLUSTERING

cluftr rebased to version 0.76.0 and fully supported

The `cluftr` packages provide a tool for transforming and analyzing cluster configuration formats. They can be used to assist with migration from an older stack configuration to a newer configuration that leverages Pacemaker. The `cluftr` tool, previously available as a Technology Preview, is now fully supported. 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>.

The `cluftr` packages have been upgraded to upstream version 0.76.0, which provides a number of bug fixes and new features. Among the notable updates are the following:

- When converting either CMAN + RGManager stack specific configuration into the respective Pacemaker configuration (or sequence of `pcs` commands) with the `ccs2pcs*` families of commands, the `cluftr` tool no longer refuses to convert entirely valid lvm resource agent configuration.
- When converting CMAN-based configuration into the analogous configuration for a Pacemaker stack with the `ccs2pcs` family of commands, some resources related configuration bits previously lost in processing (such as maximum number of failures before returning a failure to a status check) are now propagated correctly.
- When producing `pcs` commands with the `cib2pcs` and `pcs2pcscmd` families of `cluftr` commands, proper finalized syntax is now used for the alert handlers definitions, for which the (default) behavior of single-step push of the configuration changes is now respected.
- When producing `pcs` commands, the `cluftr` tool now supports a preferred ability to generate `pcs` commands that will update only the modifications made to a configuration by means of a differential update rather than a pushing a wholesale update of the entire configuration. Likewise when applicable, the `cluftr` tool now supports instructing the `pcs` tool to configure user permissions (ACLs). For this to work across the instances of various major versions of the document schemas, `cluftr` gained the notion of internal on-demand format upgrades, mirroring the internal mechanics of `pacemaker`. Similarly, `cluftr` is now capable of configuring the `bundle` feature.
- In any script-like output sequence such as that produced by the `ccs2pcscmd` and `pcs2pcscmd` families of `cluftr` commands, the intended shell interpreter is now emitted as a first, commented line as also understood directly by the operating system in order to clarify where Bash rather than a mere POSIX shell is expected. This might have been misleading under some circumstances in the past.
- The Bash completion file for `cluftr` no longer fails to work properly when the `=` character is to specify an option's value in the sequence being completed.
- The `cluftr` tool now properly detects interactive use on a terminal so as to offer more convenient representation of the outputs, and also provides better diagnostics for some previously neglected error conditions. (BZ#[1387424](#), BZ#[1381522](#), BZ#[1440876](#), BZ#[1381531](#), BZ#[1381565](#))

Support for quorum devices in a Pacemaker cluster

Red Hat Enterprise Linux 7.4 provides full support for quorum devices, previously available as a Technology Preview. This feature provides the ability to configure a separate quorum device (QDevice) which acts as a third-party arbitration device for the cluster. Its primary use is to allow a cluster to sustain more node failures than standard quorum rules allow. A quorum device is recommended for

clusters with an even number of nodes and highly recommended for two-node clusters. For information on configuring a quorum device, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/High_Availability_Add-On_Reference/. (BZ#1158805)

Support for Booth cluster ticket manager

Red Hat Enterprise Linux 7.4 provides full support for a Booth cluster ticket manager. This feature, previously available as a Technology Preview, allows you to configure multiple high availability clusters in separate sites that communicate through a distributed service to coordinate management of resources. The Booth ticket manager facilitates a consensus-based decision process for individual tickets that ensure that specified resources are run at only one site at a time, for which a ticket has been granted. For information on configuring multi-site clusters with the Booth ticket manager, see the https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/High_Availability_Add-On_Reference/ (BZ#1302087, BZ#1305049)

Support added for using shared storage with the SBD daemon

Red Hat Enterprise Linux 7.4 provides support for using the SBD (Storage-Based Death) daemon with a shared block device. This allows you to enable fencing by means of a shared block-device in addition to fencing by means of a watchdog device, which had previously been supported. The `fence-agents` package now provides the `fence_sbd` fence agent which is needed to trigger and control the actual fencing by means of an RHCS-style fence agent. SBD is not supported on Pacemaker remote nodes. (BZ#1413951)

Full support for CTDB resource agent

The CTDB resource agent used to implement a Samba deployment is now supported in Red Hat Enterprise Linux. (BZ#1077888)

The High Availability and Resilient Storage Add-Ons are now available for IBM POWER, little endian

Red Hat Enterprise Linux 7.4 adds support for the High Availability and Resilient Storage Add-Ons for the IBM POWER, little endian architecture. Note that this support is provided only for cluster nodes running on PowerVM on POWER8 servers. (BZ#1289662, BZ#1426651)

pcs now provides the ability to set up a cluster with encrypted corosync communication

The `pcs cluster setup` command now supports a new `--encryption` flag that controls the setting of corosync encryption in a cluster. This allows users to set up a cluster with encrypted corosync communication in a not entirely trusted environment. (BZ#1165821)

New commands for supporting and removing remote and guest nodes

Red Hat Enterprise Linux 7.4 provides the following new commands for creating and removing remote and guest nodes:

- `pcs cluster node add-guest`
- `pcs cluster node remove-guest`
- `pcs cluster node add-remote`
- `pcs cluster node remove-remote`

These commands replace the `pcs cluster remote-node add` and `pcs cluster remote-node remove` commands, which have been deprecated. (BZ# 1176018, BZ#1386512)

Ability to configure pcsd bind addresses

You can now configure `pcsd` bind addresses in the `/etc/sysconfig/pcsd` file. In previous releases, `pcsd` could bind to all interfaces, a situation that is not suitable for some users. By default, `pcsd` binds to all interfaces. (BZ#1373614)

New option to the `pcs resource unmanage` command to disable monitor operations

Even when a resource is in unmanaged mode, monitor operations are still run by the cluster. That may cause the cluster to report errors the user is not interested in as those errors may be expected for a particular use case when the resource is unmanaged. The `pcs resource unmanage` command now supports the `--monitor` option, which disables monitor operations when putting a resource into unmanaged mode. Additionally, the `pcs resource manage` command also supports the `--monitor` option, which enables the monitor operations when putting a resource back into managed mode. (BZ#1303969)

Support for regular expressions in `pcs` command line when configuring location constraints

`pcs` now supports regular expressions in location constraints on the command line. These constraints apply to multiple resources based on the regular expression matching resource name. This simplifies cluster management as one constraint may be used where several were needed before. (BZ#1362493)

Specifying nodes in fencing topology by a regular expression or a node attribute and its value

It is now possible to specify nodes in fencing topology by a regular expression applied on a node name and by a node attribute and its value.

For example, the following commands configure nodes `node1`, `node2`, and `node3` to use fence devices `apc1` and `apc2`, and nodes `node4`, `node5`, and `node6` to use fence devices `apc3` and `apc4`.

```
pcs stonith level add 1 "regexp%node[1-3]" apc1,apc2
pcs stonith level add 1 "regexp%node[4-6]" apc3,apc4
```

The following commands yield the same results by using node attribute matching.

```
pcs node attribute node1 rack=1
pcs node attribute node2 rack=1
pcs node attribute node3 rack=1
pcs node attribute node4 rack=2
pcs node attribute node5 rack=2
pcs node attribute node6 rack=2
pcs stonith level add 1 attrib%rack=1 apc1,apc2
pcs stonith level add 1 attrib%rack=2 apc3,apc4
```

(BZ#1261116)

Support for Oracle 11g for the resource agents `Oracle` and `OraLsnr`

Red Hat Enterprise Linux 7.4 provides support for Oracle Database 11g for the `Oracle` and `OraLsnr` resource agents used with Pacemaker. (BZ#1336847)

Support for using SBD with shared storage

Support has been added for configured SBD (Storage-Based Death) with shared storage using the `pcs` commands. For information on SBD fencing, see <https://access.redhat.com/articles/2943361>. (BZ#1413958)

Support for `NodeUtilization` resource agent

Red Hat Enterprise Linux 7.4 supports the `NodeUtilization` resource agent. The

NodeUtilization agent can detect the system parameters of available CPU, host memory availability, and hypervisor memory availability and add these parameters into the CIB. You can run the agent as a clone resource to have it automatically populate these parameters on each node. For information on the **NodeUtilization** resource agent and the resource options for this agent, run the **pcs resource describe NodeUtilization** command. For information on utilization and placement strategy in Pacemaker, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/High_Availability_Add-On_Reference/s1-utilization-HAAR.html. (BZ#1430304)

CHAPTER 7. COMPILER AND TOOLS

pcp rebased to version 3.11.8

The Performance Co-Pilot application (PCP) has been upgraded to upstream version 3.11.81, which provides a number of bug fixes and enhancements over the previous version. Notable enhancements include:

- A new client tool **pcp2influxdb** has been added to allow export of performance metric values to the **influxdb** database.
- New client tools **pcp-mpstat** and **pcp-pidstat** have been added to allow retrospective analysis of mpstat and pidstat values.
- New performance metrics have been added for device mapper, **Ceph** devices, cpusched cgroups, per-processor soft IRQs, **buddyinfo**, **zoneinfo**, shared memory, **libvirt**, same-page-sharing, **lio**, **Redis**, and **Docker**.
- Additional performance metrics from several subsystems are now available for a variety of PCP analysis tools. (BZ#[1423020](#))

systemtap rebased to version 3.1

The systemtap package has been upgraded to upstream version 3.1, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- The probes for system calls no longer default to being based on **debuginfo** information.
- Support for probing **Python** functions has been added.
- Access to **Java** function parameters has been made more uniform.
- The performance of statistical aggregate variables has been improved.
- A new statistics operator **@variance** has been added.
- More options for fetching and setting user-space values have been added.
- NFS monitoring has been improved with sample

scripts and tapset compatibility fixes. (BZ#[1398393](#), BZ#[1416204](#), BZ#1433391)

valgrind rebased to version 3.12

The valgrind package has been upgraded to upstream version 3.12, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- A new option **--ignore-range-below-sp** has been added to the **memcheck** tool to ignore memory accesses below the stack pointer. This is a generic replacement for the now deprecated option **--workaround-gcc296-bugs=yes**.
- The maximum number of callers in a suppression entry generated by the **--gen-suppressions=yes** option is now equal to the value given by the **--num-callers** option.
- The cost of instrumenting code blocks for the most common use case, the **memcheck** tool on the AMD64 and Intel 64 architectures, has been reduced.
- Performance has been improved for debugging programs which discard a lot of instruction address ranges of 8KB or less.

- Support for IBM Power 9 (ISA 3.0) architecture has been added.
- Partial support for AMD FMA4 instructions has been added.
- Support for cryptographic and CRC instructions on the 64-bit ARM architecture version 8 has been added. (BZ#1391217)

New package: unitsofmeasurement

The unitsofmeasurement package enables expressing units of measurement in Java code. With the new API for units of measurement, handling of physical quantities is easier and less error-prone. The package's API is efficient in use of memory and resources. (BZ#1422263)

SSL/TLS certificate verification for HTTP clients is now enabled by default in the Python standard library

The default global setting for HTTP clients has been changed in the Python standard library to verify SSL/TLS certificates by default. Customers using the file-based configuration are not affected. For details, see <https://access.redhat.com/articles/2039753>. (BZ#1219110)

Support for %gemspec_add_dep and %gemspec_remove_dep has been added

This update adds support for the %gemspec_add_dep and %gemspec_remove_dep macros. These macros allow easier adjustment of rubygem-* package dependencies. In addition, all current macros have been extended to improve support for pre-release version of packages. (BZ#1397390)

ipmitool rebased to version 1.8.18

The ipmitool package has been upgraded to upstream version 1.8.18, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- The PEF user interface has been redesigned
- A new subcommand `lan6` has been added for IP version 6 local area network parameters
- Support for VITA-specific sensor types and events has been added
- Support for HMAC_MD5 and HMAC_SHA256 encryption has been added
- Support for checking PICMG extension 5.x has been added
- Support for USB medium as a new communication interface has been added
- The USB driver has been enabled by default for GNU Linux systems (BZ#1398658)

Ishw updated for the little-endian variant of IBM Power

The Ishw packages, which provide detailed information on the hardware configuration of a machine, have been updated for the little-endian variant of IBM Power System. (BZ#1368704)

perf now supports uncore events on Intel Xeon v5

With this update, Performance analysis tool for Linux (perf) has been updated to support uncore events on Intel Xeon v5 server CPU. These events provide additional performance monitoring information for advanced users. (BZ#1355919)

dmidecode updated

The dmidecode package has been updated to a later version, which provides several bug fixes and hardware enablement improvements. (BZ#1385884)

iSCSI now supports configuring the ALUA operation by using targetcli

With multiple paths from the initiator to a target, you can use Asymmetric Logical Unit Assignment

(ALUA) to configure preferences for how to use the paths in a non-uniform, preferential way. The Linux-IO (LIO) kernel target has always supported this feature. With this update, you can use the `targetcli` command shell to configure the ALUA operation. (BZ#1243410)

jansson rebased to version 2.10

The `jansson` library has been updated to version 2.10, which provides several bug fixes and enhancements over the previous version. Notably, interfaces have been added to support the `clevis`, `tang` and `jose` applications. (BZ#1389805)

A new compatibility environmental variable for `egrep` and `fgrep`

In an earlier `grep` rebase, the `egrep` and `fgrep` commands were replaced by `grep -E` and `grep -F` respectively. This change could affect customer scripts because only `grep` was shown in the output of the `ps` command. To prevent such problems, this update introduces a new compatibility environmental variable: `GREP_LEGACY_EGREP_FGREP_PS`. To preserve showing `egrep` and `fgrep` in `ps` output, set the variable to 1:

```
GREP_LEGACY_EGREP_FGREP_PS=1
```

(BZ#1297441)

`lastcomm` now supports the `--pid` option

The `lastcomm` command now supports the `--pid` option. This option shows the process ID (PID) and parent-process ID (PPID) for each record if supported by the kernel. (BZ#1255183)

New package: `perl-Perl4-CoreLibs`

A new `perl-Perl4-CoreLibs` package is now available in the Base channel of Red Hat Enterprise Linux 7. This package contains libraries that were previously available in Perl 4 but were removed from Perl 5.16, which is distributed with Red Hat Enterprise Linux 7. In the previous release, these libraries were provided in a Perl subpackage through the Optional channel. (BZ#1366724)

`tar` now follows symlinks to directories when extracting from the archive

This update adds the `--keep-directory-symlink` option to the `tar` command. This option changes the behavior of `tar` when it encounters a symlink with the same name as the directory that it is about to extract. By default, `tar` would first remove the symlink and then proceed extracting the directory. The `--keep-directory-symlink` option disables this behavior and instructs `tar` to follow symlinks to directories when extracting from the archive. (BZ#1350640)

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

The `Net::SSLeay` Perl module has been updated to support explicit specification of the TLS protocol versions 1.1 or 1.2 to improve security, and the `IO::Socket::SSL` module has been updated accordingly. When a new `IO::Socket::SSL` object is created, it is now possible to restrict the TLS version to 1.1 or 1.2 by setting the `SSL_version` option to `TLSv1_1` or `TLSv1_2` respectively. Alternatively, `TLSv11` and `TLSv12` can be used. Note that these values are case-sensitive. (BZ#1335035)

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

The `Net::SSLeay` Perl module has been updated to support explicit specification of the TLS protocol version, which can be used for improving security. To restrict TLS version to 1.1 or 1.2, set the `Net::SSLeay::ssl_version` variable to `11` or `12` respectively. (BZ#1335028)

`wget` now supports specification of the TLS protocol version

Previously, the `wget` utility used the highest TLS protocol version 1.2 by default when connecting to a remote server. With this update, `wget` has been enhanced to allow the user to explicitly select the TLS

protocol minor version by adding the `--secure-protocol=TLSv1_1` or `--secure-protocol=TLSv1_2` command-line options to the `wget` command. (BZ#1439811)

tcpdump rebased to version 4.9.0

The `tcpdump` package has been upgraded to upstream version 4.9.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Many security vulnerabilities have been fixed
- Numerous improvements have been made in the dissection of popular network protocols
- The default `snaplen` feature has been increased to 262144 bytes
- The capture buffer has been enlarged to 4 MiB (BZ#1422473)

The option to set capture direction for tcpdump changed from -P to -Q

Previously, the `tcpdump` utility in Red Hat Enterprise Linux used the `-P` option to set the capture direction, while the upstream version used `-Q`. The `-Q` option has been implemented and is now preferred. The `-P` option retains the previous function as an alias of `-Q`, but displays a warning. (BZ#1292056)

OpenJDK now supports SystemTap on the 64-bit ARM architecture

The OpenJDK platform now supports introspection with the `SystemTap` instrumentation tool on the 64-bit ARM architecture. (BZ#1373986)

sos rebased to version 3.4

The `sos` package has been updated to upstream version 3.4, which provides a number of enhancements, new features, and bug fixes, including:

- New plug-ins have been added for `ceph_ansible`, `collectd`, `crypto`, `dracut`, `gnocchi`, `jars`, `nfsganesha`, `nodejs`, `npm`, `openstack_ansible`, `openstack_instack`, `openstack_manila`, `salt`, `saltmaster`, and `storageconsole`
- API plug-in enhancements
- Internationalisation updates
- The networking plug-in no longer crashes when a network name contains the single quote character `'`
- The `foreman-debug` plug-in is now run with a longer timeout to prevent incomplete `foreman-debug` information collected
- Certain private SSL certificate files are no longer collected (BZ#1414879)

targetd rebased to version 0.8.6

The `targetd` packages have been upgraded to upstream version 0.8.6, which provides a number of bug fixes and enhancements over the previous version. Notably, the `targetd` service now runs on either Python 2 or Python 3 run time, and the following APIs have been added: `initiator_list`, `access_group_list`, `access_group_create`, `access_group_destroy`, `access_group_init_add`, `access_group_init_del`, `access_group_map_list`, `access_group_map_create`, and `access_group_map_destroy`.

Notable bug fixes include:

- **targetd** is now compliant with JSON-RPC response version 2.0.
- the **export_create** API can now be used to map the same LUN to multiple initiators.
- **targetd** now ensures that SSL certificates are present at start-up. (BZ# [1162381](#))

shim rebased to version 12-1

With this update, the shim package has been upgraded to upstream version 12-1, which provides a number of bug fixes and enhancements over the previous version. Notably, the support for 32-bit UEFI firmware and Extensible Firmware Interface (EFI) utilities has been added. (BZ#1310766)

rubygem-abrt rebased to version 0.3.0

The rubygem-abrt package has been rebased to version 0.3.0, which provides several bug fixes and enhancements over the previous version. Notably:

- The **Ruby ABRT** handler now supports **uReports**, automatic anonymous microreports. With **uReports** enabled, developers are promptly notified about application issues and are able to fix bugs and resolve problems faster.
- Previously, when a **Ruby** application was using **Bundler** to manage its dependencies and an error occurred, an incorrect logic was used to load components of the **Ruby ABRT** handler. Consequently, an unexpected **LoadReport** error was reported to the user instead of a proper **ABRT** report. The loading logic has been fixed, and the **Ruby** application errors are now correctly handled and reported using **ABRT**. (BZ#[1418750](#))

New package: http-parser

The new **http-parser** package provides a utility for parsing HTTP messages. It parses both requests and responses. The parser is designed to be used in applications managing HTTP performance. It does not make any syscalls or allocations, it does not buffer data, and it can be interrupted at any time. Depending on your architecture, it only requires about 40 bytes of data per message stream. (BZ#1393819)

Intel and IBM POWER transactional memory support for all default POSIX mutexes

The default POSIX mutexes can be transparently substituted with Intel and IBM POWER transactional memory support, which significantly reduces the cost of lock acquisition. To enable transactional memory support for all default POSIX mutexes, set the **RHEL_GLIBC_TUNABLES=glibc.elision.enable** environment variable to **1**. As a result, performance of some applications can be improved.

Developers are advised to use profiling to decide whether enabling of this feature improves performance for their applications. (BZ#[841653](#), BZ#731835)

glibc now supports group merging

The ability to merge group members from different Name Service modules has been added to **glibc**. As a result, management of centralized user access control and group membership across multiple hosts is now easier. (BZ#[1298975](#))

glibc now supports optimized string comparison functions on The IBM POWER9 architecture

The string comparison functions **strcmp** and **strncmp** from the **glibc** library have been optimized for the IBM POWER9 architecture. (BZ#1320947)

Improved performance for dynamically loaded libraries using the Intel SSE, AVX and AVX512 features

Dynamic library loading has been updated for libraries using the Intel SSE, AVX, and AVX512 features. As a result, performance while loading these libraries has improved. Additionally, support for LD_AUDIT-style auditing has been added. (BZ#[1421155](#))

elfutils rebased to version 0.168

The elfutils package has been upgraded to upstream version 0.168, which provides a number of bug fixes and enhancements:

- The option `--symbols` of the `eu-readelf` utility now allows selecting the section for displaying symbols.
- New functions for the creation of ELF/DWARF string tables have been added to the `libdw` library.
- The `DW_LANG_PL1` constant has been changed to `DW_LANG_PLI`. The previous name is still accepted.
- The return type of the `gelf_newehdr` and `gelf_newphdr` functions from the `libelf` library has been changed to `void*` for source compatibility with other `libelf` implementations. This change retains binary compatibility on all platforms supported by Red Hat Enterprise Linux. (BZ#[1400302](#))

bison rebased to version 3.0.4

The bison package has been upgraded to upstream version 3.0.4, which provides a number of bug fixes and enhancements:

- Endless diagnostics caused by caret errors have been fixed.
- The `-Werror=CATEGORY` option has been added to treat specified warnings as errors. The warnings do not have to be explicitly activated using the `-W` option.
- Many improvements in handling of precedence rules and useless rules.

Additionally, the following changes breaking backward compatibility have been introduced:

- The following features have been deprecated: `YYFAIL`, `YYLEX_PARAM`, `YYPARSE_PARAM`, `yystype`, `yytype`
- Missing semicolons at the end of actions are no longer automatically added.
- To use Bison extensions with the `autoconf` utility versions 2.69 and earlier, pass the option `-Wno-yacc` to `(AM_)YFLAGS`. (BZ#[1306000](#))

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

Previously, when connecting to a new system via TLS/SSL, the `Mut t` email client required the user to save the certificate. With this update, the system Certificate Authority (CA) bundle is set in `Mut t` by default. As a result, `Mut t` now connects via SSL/TLS to hosts with a valid certificate without prompting the user to approve or reject the certificate. (BZ#[1388511](#))

objdump mixed listing speed up

Previously, the BFD library for parsing DWARF debug information and locating source code was very slow. The BFD library is used by the `objdump` tool. As a consequence, `objdump` became significantly slower when producing a mixed listing of source code and disassembly. Performance of the BFD library has been improved. As a result, producing a mixed listing with `objdump` is faster. (BZ#[1366052](#))

ethtool support for human readable output from the fjes driver

The **ethtool** utility has been enhanced to provide a human readable form of register dump output from the **fjes** driver. As a result, users of **ethtool** can inspect the Fujitsu Extended Socket Network Device driver more comfortably. (BZ#1402701)

ecj rebased to version 4.5.2

The **ecj** package has been upgraded to upstream version 4.5.2, which provides a number of bug fixes and enhancements over the previous version. Notably, support for features added to the Java language in version 8 has been completed. As a result, compilation of Java code using Java 8 features no longer fails. This includes cases where code not using Java 8 features referenced code using these features, such as system classes provided by the Java Runtime Environment. (BZ#[1379855](#))

rhino rebased to version 1.7R5

The **rhino** package has been upgraded to upstream version 1.7R5, which provides a number of bug fixes and enhancements over the previous version. Notably, the former problem with an infinite loop while parsing regular expressions has been fixed. Applications using **Rhino** that previously encountered this bug now function correctly. (BZ#[1350331](#))

scap-security-guide and oscap-docker now support containers

The user can now use the **oscap-docker** utility and the SCAP Security Guide to assess compliance of container or container image without encountering false positive results. Tests that make no sense in container context, such as partitioning, has been set to the **not applicable** value, and containers can be now scanned with a selected security policy. (BZ#[1404392](#))

CHAPTER 8. DESKTOP

GNOME rebased to version 3.22.3

The GNOME Desktop has been updated to upstream version 3.22.3, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Overhauled desktop notifications
- Built-in integration with world clocks and media players
- Automatically adjusting screen brightness (for systems with an integrated light sensor)
- In many applications, support for a standard dialog to document keyboard shortcuts
- Improvements to several setting panels (printer, mouse, touchpad, keyboard shortcuts)
- Option to rename multiple files at once
- Built-in support for compressed files and for Google Drive
- Undo support for trash (BZ#[1383353](#))

The `xorg-x11-drv-libinput` driver has been added to the X.Org input drivers

The `xorg-x11-drv-libinput` X.Org driver is a wrapper driver for the low-level `libinput` library. This update adds the driver to the X.Org input drivers. After you install `xorg-x11-drv-libinput`, it is possible to remove the `xorg-x11-drv-synaptics` driver and get access to some of the improved input device handling offered by `libinput`. (BZ#[1413811](#))

Change of default driver for some Intel and nVidia Hardware

This change affects:

- 4th Generation Intel Core Processors and later
- nVidia GeForce 8 hardware and later

The default DDX driver has changed to be `xf86-video-modesetting`.

Previously the defaults were `xf86-video-nouveau` and `xf86-video-intel` for nVidia and Intel hardware respectively. (BZ#[1404868](#))

`dconf-editor` is now provided by a separate package

The upstream `dconf` team have split the `dconf-editor` into its own package. This release reflects that change.

In addition, the user interface has been redesigned in version 3.22.

- The tree view on the left side has been removed.
- Keys and directories are now shown in the same window.
- The ability to go back in the hierarchy has moved to the path shown in the header-bar. (BZ#[1388931](#))

CHAPTER 9. FILE SYSTEMS

SELinux security labels are now supported on the OverlayFS file system

With this update, the OverlayFS file system now supports SELinux security labels. When using Docker containers with the OverlayFS storage driver, you no longer have to configure Docker to disable SELinux support for the containers. (BZ#[1297929](#))

NFSv4 server is now fully supported

NFS over RDMA (NFSv4) server, previously provided as a Technology Preview, is now fully supported when accessed by Red Hat Enterprise Linux clients. For more information on NFSv4 see the following section in the Red Hat Enterprise Linux 7 Storage Administration Guide: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html-single/Storage_Administration_Guide/index.html#nfs-rdma (BZ#1400501)

autofs now supports the browse options of amd format maps

The browse functionality of sun format maps makes available automount points visible in directory listings of mounted automount-managed mounts and is now also available for `autofs` `amd` format maps.

You can now add mount point sections to the `autofs` configuration for `amd` format mounts, in the same way automount points are configured in `amd`, without the need to also add a corresponding entry to the master map. As a result, you can avoid having incompatible master map entries in the `autofs` master map within shared multi-vendor environments.

The `browsable_dirs` option can be used in either the `autofs [ amd ]` configuration section, or following `amd` mount point sections. The `browsable` and `utimeout` map options of `amd` type `auto` map entries can also be used.

Note that the `browsable_dirs` option can be set only to `yes` or `no`. (BZ#[1367576](#))

To make searching logs easier, autofs now provides identifiers of mount request log entries

For busy sites, it can be difficult to identify log entries for specific mount attempts when examining mount problems. The entries are often mixed with other concurrent mount requests and activities if the log recorded a lot of activity. Now, you can quickly filter entries for specific mount requests if you enable adding a mount request log identifier to mount request log entries in the `autofs` configuration. The new logging is turned off by default and is controlled by the `use_mount_request_log_id` option, as described in the `autofs.conf` file. (BZ#1382093)

GFS2 on IBM z Systems is now supported in SSI environments

Starting with Red Hat Enterprise Linux 7.4, GFS2 on IBM z Systems (Resilient Storage on the `s390x` add-on) is supported in z/VM Single System Image (SSI) environments, with multiple central electronics complexes (CECs). This allows the cluster to stay up even when logical partitions (LPARs) or CECs are restarted. Live migration is not supported due to the real-time requirements of High Availability (HA) clustering. The maximum node limit of 4 nodes on IBM z Systems still applies. For information on configuring high availability and resilient storage for IBM z systems, see <https://access.redhat.com/articles/1543363>. (BZ#1273401)

gfs2-utils rebased to version 3.1.10

The `gfs2-utils` packages have been upgraded to upstream version 3.1.10, which provides a number of bug fixes and enhancements over the previous version. Notably, this update provides:

- various checking and performance improvements of the `fsck.gfs2` command

- better handling of odd block device geometry in the `mkfs.gfs2` command.
- `gfs2_edit savemeta` leaf chain block handling bug fixes.
- handling UUIDs by the `libuuid` library instead of custom functions.
- new `--enable-gprof` configuration option for profiling.
- documentation improvements. (BZ#1413684)

FUSE now supports `SEEK_HOLE` and `SEEK_DATA` in `lseek` calls

This update provides the `SEEK_HOLE` and `SEEK_DATA` features for the Filesystem in Userspace (FUSE) `lseek` system call. Now, you can use FUSE `lseek` to adjust the offset of the file to the next location in the file that contains data, with `SEEK_DATA`, or a hole, with `SEEK_HOLE`. (BZ#1306396)

NFS server now supports limited copy-offload

The NFS server-side copy feature now allows the NFS client to copy file data between two files that reside on the same file system on the same NFS server without the need to transmit data back and forth over the network through the NFS client. Note that the NFS protocol also allows copies between different file systems or servers, but the Red Hat Enterprise Linux implementation currently does not support such operations. (BZ#1356122)

SELinux is supported for use with GFS2 file systems

Security Enhanced Linux (SELinux) is now supported for use with GFS2 file systems. Since use of SELinux with GFS2 incurs a small performance penalty, you may choose not to use SELinux with GFS2 even on a system with SELinux in enforcing mode. For information on how to configure this, see https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Global_File_System_2/index.html. (BZ#437984)

NFSv4.1 client and server now support Kerberos authentication

This update adds Kerberos authentication support for NFS over RDMA (NFSv4.1 client and server) to allow you to use `krb5`, `krb5i`, and `krb5p` authentication with NFSv4.1 features. You can now use Kerberos with NFSv4.1 for secure authentication of each Remote Procedure Call (RPC) transaction. Note that you need version 1.3.0-0.36 or higher of the `nfs-utils` package to be installed to use Kerberos with NFSv4.1. (BZ#1401797)

`rpc.idmapd` now supports obtaining NFSv4 ID Domains from DNS

The NFS domain name that is used in the ID mapping can now be retrieved from DNS. If the `Domain` variable is not set in the `/etc/idmapd.conf` file, DNS is queried to search for the `_nfsv4idmapdomain` Text record. If a value is found, it is used as the NFS domain. (BZ#980925)

NFSv4.1 is now the default NFS mount protocol

Prior to this update, NFSv4.0 was the default NFS mount protocol. NFSv4.1 provides significant feature improvements over NFSv4.0, such as sessions, pNFS, parallel OPENS, and session trunking. With this update, NFSv4.1 is the default NFS mount protocol.

If you have already specified the mount protocol minor version, this update causes no change in behavior. This update causes a change in behavior if you have specified NFSv4 without a specific minor version, provided the server supports NFSv4.1. If the server only supports NFSv4.0, the mount remains a NFSv4.0 mount. You can retain the original behavior by specifying `0` as the minor version:

- on the mount command line,
- in the `/etc/fstab` file,

- or in the `/etc/nfsmount.conf` file. (BZ#[1375259](#))

Setting `nfs-utils` configuration options has been centralized in `nfs.conf`

With this update, `nfs-utils` uses configuration centralized in the `nfs.conf` file, which is structured into stanzas for each `nfs-utils` program. Each `nfs-utils` program can read the configuration directly from the file, so you no longer need to use the `systemctl restart nfs-config.service` command, but restart only the specific program. For more information, see the `nfs.conf(5)` manual page.

For compatibility with earlier releases, the older `/etc/sysconfig/nfs` configuration method is still available. However, it is recommended to avoid specifying configuration settings in both the `/etc/sysconfig/nfs` and `/etc/nfs.conf` file. (BZ#[1418041](#))

Locking performance for NFSv4.1 mounts has been improved for certain workloads

NFSv4 clients poll the server at an interval to obtain a lock under contention. As a result, the locking performance for contented locks for NFSv4 is slower than the performance of NFSv3.

The `CB_NOTIFY_LOCK` operation has been added to the NFS client and server, so NFSv4.1 and later allow servers to call back to clients waiting on a lock.

This update improves the locking performance for contented locks on NFSv4.1 mounts for certain workloads. Note that the performance might not improve for longer lock contention times. (BZ#[1377710](#))

CHAPTER 10. HARDWARE ENABLEMENT

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#1386133)

New Wacom driver introduced in 7.4 to support upcoming tablets

With this update, a new Wacom driver was introduced to support recently-released and upcoming tablets meanwhile the current driver continues supporting previously-released tablets.

Notable features:

- Wacom 27QHT (touch) is now supported
- ExpressKey Remote (BZ#1385026)

Wacom kernel driver now supports ThinkPad X1 Yoga touch screen

With this update, the support for the ThinkPad X1 Yoga touch screen has been added to the Wacom kernel driver. As a result, the touch screen can be properly used when running Red Hat Enterprise Linux 7 on these machines. (BZ#1388646)

The touch functionality has been added to the Wacom Cintiq 27 QHDT tablets

This update adds support of the touch functionality for the Wacom Cintiq 27 QHDT tablets, which makes it possible to properly use the touch screen when running Red Hat Enterprise Linux 7 on these machines. (BZ#1391668)

AMDGPU now supports the Southern Islands, Sea Islands, Volcanic Islands and Arctic Islands chipsets

The support for the Southern Islands, Sea Islands, Volcanic Islands and Arctic Islands chipsets has been added. The AMDGPU graphics driver is the next generation family of open source graphics drivers for the latest AMD/ATI Radeon graphics cards. It is based on the Southern Islands, Sea Islands, Volcanic Islands, and Arctic Islands chipsets. It is necessary to install the proper firmware or microcode for the card provided by the `linux-firmware` package. (BZ#1385757)

Support added for the AMD mobile graphics

Support for the AMD mobile graphics based on the Polaris architecture has been added. The Polaris architecture is based on the Arctic Islands chipsets. It is necessary to install the proper firmware or microcode for the card provided by the `linux-firmware` package. (BZ#1339127)

Netronome NFP devices are supported

With this update, the `nfp` driver has been added to the Linux kernel. As a result, the Netronome Network Flow Processor (Netronome NFP 4000/6000 VF) devices are now supported on Red Hat Enterprise Linux 7. (BZ#1377767)

nvme-cli rebased to version 1.3

The `nvme-cli` utility has been updated to version 1.3, which includes support for Nonvolatile Memory Express (NVMe). With the support for NVMe, you can find the targets over Remote Direct Memory Access (RDMA) and connect to these targets. (BZ#1382119)

The queued spinlocks have been implemented into the Linux kernel

This update has changed the spinlock implementation in the kernel from ticket spinlocks to queued

spinlocks on AMD64 and Intel 64 architectures. The queued spinlocks are more scalable than the ticket spinlocks. As a result, the system performance has been improved especially on Symmetric Multi Processing (SMP) systems with large number of CPUs. The performance now increases more linearly with increasing number of the CPUs. Note that because of this change in the spinlock implementation, kernel modules built on Red Hat Enterprise Linux 7 might not be loadable on kernels from earlier releases. Kernel modules released in Red Hat Enterprise Linux (RHEL) versions earlier than 7.4 are loadable on the kernel released in RHEL 7.4. (BZ#1241990)

rapl now supports Intel Xeon v2 servers

The Intel `rapl` driver has been updated to support Intel Xeon v2 servers. (BZ#1379590)

Further support for Intel Platform Controller Hub [PCH] devices

The kernel has been updated to enable support for new Intel PCH hardware on the Intel Xeon Processor E3 v6 Family CPUs. (BZ#1391219)

Included genwqe-tools to enable use of hardware accelerated zLib on IBM Power and s390x

The `genwqe-tools` package allows users of IBM Power and s390x hardware to utilize FPGA based PCIe cards for zLib compression and decompression processes.

These tools enable use of RFC1950, RFC1951 and RFC1952 compliant hardware to increase performance. (BZ#1275663)

librtas rebased to version 2.0.1

The `librtas` packages have been upgraded to upstream version 2.0.1, which provides a number of bug fixes and enhancements over the previous version. Notably, this update changes the soname of the provided libraries: `librtas.so.1` changes to `librtas.so.2`, and `librtasevent.so.1` changes to `librtasevent.so.2`. (BZ#1380656)

The NFP driver

The Network Flow Processor (NFP) driver has been backported from version 4.11 of the Linux kernel. This driver supports Netronome NFP4000 and NFP6000 based cards working as an advanced Ethernet NIC. The driver works with both SR-IOV physical and virtual functions. (BZ#1406197)

Enable latest NVIDIA cards in Nouveau

This update includes enablement code to ensure that higher-end NVIDIA cards based on the Pascal platform work correctly. (BZ#1330457)

Support for Wacom ExpressKey Remote

The Wacom ExpressKey Remote (EKR) is now supported in Red Hat Enterprise Linux 7. EKR is an external device that allows you to access shortcuts, menus and commands. (BZ#1346348)

Wacom Cintiq 27 QHD now supports ExpressKey Remote

With this update, Wacom Cintiq 27 QHD tablets support ExpressKey Remote (EKR). EKR is an external device that allows you to access shortcuts, menus and commands. (BZ#1342990)

CHAPTER 11. INSTALLATION AND BOOTING

Anaconda enables users to set RAID chunk size

This update allows the user to set the `--chunksize` parameter for the `raid` utility in a kickstart file to specify the chunk size of a RAID storage, in KiB. Using the `--chunksize` parameter overrides the default one. As a result, the new chunk size can prevent a negative performance impact the default value might have. (BZ#1332316)

Anaconda text mode now supports IPoIB interfaces

This update adds support for IP over InfiniBand (IPoIB) network interfaces during manual installation in text mode. You can now view IPoIB interface status information and change interface configuration. (BZ#1366935)

`inst.debug` enables a more convenient debugging of Anaconda installation issues

This update adds the ability to save logs related to the initial state of the machine by starting the Anaconda installation with the `inst.debug` boot option. This option stores three additional logs, `lsblk`, `dmesg` and `lvm dump`, in the `/tmp/pre-anaconda-logs/` directory, allowing a more convenient debugging of issues which occurred during the installation. (BZ#1255659)

Kickstart installation failure automatically triggers `%onerror` scripts

This enhancement makes sure that the `%onerror` sections in a kickstart file are run if the Anaconda installation fails. The scripts can be used to collect logs automatically for further examination. As a result of this update, when a traceback or another fatal error occurs during the installation, the installer performs the `%onerror` scripts and the `%traceback` scripts check if the error was caused by a traceback. (BZ#1412538)

Anaconda can now wait for network to become available before starting the installation

In some environments, the first DHCP request can be expected to fail. Previously, the first DHCP failure caused Anaconda to proceed with the installation, which could cause problems especially with automatic installations where a connection could not be set up manually later. This update introduces a new Anaconda boot option, `inst.waitfornet=X`, which forces the installer to spend X seconds waiting for network connectivity before proceeding. The installation will continue once a connection is established, or after the specified time interval has passed. (BZ#1315160)

Multiple network locations of stage2 or Kickstart files can be specified to prevent installation failure

This update enables the specification of multiple `inst.stage2` and `inst.ks` boot options with network locations of stage2 and a Kickstart file. This avoids the situation in which the requested files cannot be reached and the installation fails because the contacted server with the stage2 or the Kickstart file is inaccessible.

With the new update, the installation failure can be avoided if multiple locations are specified. If all the defined locations are URLs, namely `HTTP`, `HTTPS`, or `FTP`, they will be tried sequentially until the requested file is fetched successfully. If there is a location that is not a URL, only the last specified location is tried. The remaining locations are ignored. (BZ#1391724)

`autopart --nohome` in a kickstart file disables the creation of `/home/` in automatic partitioning

This update adds the `--nohome` option to the `autopart` command in a kickstart file to disable automatic creation of the `/home/` partition. This enhancement avoids the need to perform manual partitioning if the `/home/` partition is to be averted. As a result of the update, the `/home` partition is not created if partitioning is done automatically. (BZ#663099)

Loading driver disks from hard disk drives and USBs enabled

This update enables loading driver disks from a hard disk drive or a similar device instead of loading them over the network or from `initrd`. The installation can proceed either using the kickstart or the boot options.

The procedure is as follows:

1. Load the driver disk on a hard disk drive, a USB or any similar device.
2. Set the label, for example, **DD**, to this device.

Notes:

For kickstart installation, add

```
driverdisk LABEL=DD:/e1000.rpm
```

to your kickstart file.

For the boot option, start the installation with

```
inst.dd=hd:LABEL=DD:/dd.rpm
```

as the boot argument.

In both the kickstart and the boot options, replace **DD** with a specific label and replace `dd.rpm` with a specific name. Use anything supported by the `inst.repo` command instead of **LABEL** to specify your hard disk drive. Do not use non-alphanumeric characters in the argument specifying the **LABEL** of the kickstart `driverdisk` command. (BZ#1377233)

Changes in automatic partitioning behavior for LVM thin pools

Previously, every Logical Volume Management (LVM) thin pool created or used in the installation, whether using Kickstart or an interactive installation, 20 % of its size reserved.

This update brings the following changes:

- If you create a LVM thin pool with automatic partitioning, 20 % of the volume group size is reserved, with a minimum of 1 GiB and a maximum of 100 GiB.
- If you use the `logvol --thinpool --grow` command in a Kickstart file, the thin pool will grow to the maximum possible size, which means no space will be left for it in the volume group to grow. In this case, you can use the `volgroup --reserved-space` or `volgroup --reserved-percent` command to leave some space in the volume group reserved, which is recommended. (BZ#1131247)

32-bit boot loaders can now boot 64-bit kernels on UEFI

This update enables booting 64-bit kernels using 32-bit boot loaders, such as `grub2-i386-efi`, on systems with UEFI firmware. (BZ#1310775)

Lorax can now ignore SSL errors

Previously, the `lorax` tool could not use HTTPS repositories with self-signed certificates. An attempt to do so resulted in an error with no way to continue. This update adds the `--noverifyssl` command line option to the utility, which can be used to skip verifying the server certificate and bypass the error. (BZ#1430483)

shim-signed rebased to version 12

With this update, the shim-signed package has been upgraded to upstream version 12, which provides a number of bug fixes and enhancements over the previous version. Notably, support for 32-bit UEFI firmware and Extensible Firmware Interface (EFI) utilities has been added. (BZ#1310764)

gnu-efi rebased to version 3.0.5.-9

With this update, the gnu-efi package has been upgraded to upstream version 3.0.5.-9, which provides a number of bug fixes and enhancements over the previous version. Notably, the support for 32-bit UEFI firmware and Extensible Firmware Interface (EFI) utilities has been added. (BZ#1310782)

Backward compatibility enabled for `killproc()` and `status()`

Prior to this update, the `/etc/rc.d/init.d/functions` script shipped in Red Hat Enterprise Linux 7 lacked some of the features of the Red Hat Enterprise Linux 6 counterpart. The `initscripts` package has been updated to add support for the `-b` option to the `killproc()` and `status()` functions in the `/etc/rc.d/init.d/functions` file. This addition enables backward compatibility for Red Hat Enterprise Linux 6 and prevents possible regressions when performing an upgrade from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7. (BZ#1428935)

DHCP_FQDN allows specifying a fully qualified domain name of the system

Previously, the `ifcfg` interface configuration files required that the `DHCP_HOSTNAME` directive was used to specify the hostname of the system. The new `initscripts` `DHCP_FQDN` directive now also allows specifying the fully qualified domain name of the system. This is a complement to the `DHCP_HOSTNAME` directive. In case both `DHCP_HOSTNAME` and `DHCP_FQDN` are specified, only `DHCP_FQDN` is used. (BZ#1260552)

You can now create thin logical volume snapshots during the installation process

This update adds support for a new Kickstart command, `snapshot`. This command allows you to create a LVM thin volume snapshot either before or after the installation. Available options are:

- `<vg_name>/<lv_name>` Specify the name of the volume group and logical volume to make a snapshot of.
- `--name=` Specify a name for the snapshot.
- `--when=` Specify `pre-install` if you want to take a snapshot before the installation begins, which can be useful if you want to preserve the state of a system before an upgrade. Alternatively, specify `post-install` to take a snapshot of a newly installed system before any additional changes are made to it.

All three options are mandatory. Also note that you can use this command multiple times in a single Kickstart file to take a snapshot both before and after the installation, or to take snapshots of multiple logical volumes. Make sure that each `--name=` parameter specifies a unique name when doing so. (BZ#1113207)

CHAPTER 12. KERNEL

The NVMe driver rebased to kernel version 4.10

The NVM-Express kernel driver has been updated to upstream kernel version 4.10, which provides a number of bug fixes and enhancements over the previous version. The most notable change is: the initial NVMe-over-Fabrics transport implementation, which uses existing RDMA NICs (Infiniband, RoCE, iWARP) and existing NVMe SSDs, has been added to the driver, but does not include support for DIF/DIX and multipathing. (BZ#1383834)

crash rebased to version 7.1.9

With this update, the crash packages have been upgraded to upstream version 7.1.9, which provides a number of bug fixes and enhancements over the previous version. (BZ#1393534)

crash now analyzes vmcore dumps for IBM Power ISA 3.0

The **crash** utility has been updated to correspond with changes in the kernel page table related to IBM Power ISA version 3.0 architecture. As a result, the **crash** utility is now able to analyze vmcore dumps of kernels on IBM Power ISA 3.0 systems. (BZ#1368711)

crash updated for IBM Power and for the little-endian variant of IBM Power

The crash packages have been updated to support IBM Power Systems and the little-endian variant of IBM Power Systems. These packages provide the core analysis suite, which is a self-contained tool that can be used to investigate live systems, as well as kernel core dumps created by the kexec-tools packages or the Red Hat Enterprise Linux kernel. (BZ#1384944)

memkind updated to version 1.3.0

The memkind library has been updated to version 1.3.0, which provides several bug fixes and enhancements over the previous version.

Notable changes include:

- A logging mechanism has been introduced.
- Hardware Locality (hwloc) has been integrated, and can be turned on using the `--with-hwloc` option.
- The symbols exposed by libmemkind.so have been cleaned up. For example, `libnuma` and `jemalloc` are no longer exposed.
- AutoHBW files have been moved to the `/memkind/autohbw/` directory, code has been refactored and tests have been added to appropriate scenarios.
- Flags improving security have been added to memkind. The flags can be turned off with the `--disable-secure` configure time option.
- The configuration of jemalloc has been changed to turn off unused features.
- Several symbols have been deprecated. For details, see the Deprecated Functionality part. (BZ#1384549)

Jitter Entropy RNG added to the kernel

This update adds the Jitter Entropy Random Number Generator (RNG), which collects entropy through CPU timing differences to the Linux kernel. This RNG is by default available through the `algif_rng` interface. The generated numbers can be added back to the kernel through the `/dev/random` file, which makes these numbers available to other `/dev/random` users. As a result, the operating system now has more sources of entropy available. (BZ#1270982)

/dev/random now shows notifications and warnings for the urandom pool**initialization**

With this update, the random driver (/dev/random), has been modified to print a message when the nonblocking pool (used by /dev/urandom) is initialized. (BZ#1298643)

fjes updated to version 1.2

The fjes driver has been updated to version 1.2, which includes a number of bug fixes and enhancements over the previous version. (BZ#1388716)

Full support for user name spaces

User name spaces (userns) that were introduced in Red Hat Enterprise Linux 7.2 as Technology Preview are now fully supported. This feature provides additional security to servers running Linux containers by improving isolation between the host and the containers. Administrators of containers are no longer able to perform administrative operations on the host, which increases security.

The default value of `user.max_user_namespaces` is 0. You can set it to a non-zero value, which stops the applications that malfunction. It is recommended that `user.max_usernamespaces` is set to a large value, such as 15000, so that the value does not need to be revisited in the normal course of operation. (BZ#1340238)

makedumpfile updated to version 1.6.1

The makedumpfile package has been upgraded to upstream version 1.6.1 as part of the kexec-tools 2.0.14 rpm, which provides a number of bug fixes and enhancements over the previous version. (BZ#1384945)

Intel Cache Allocation Technology is supported

This update adds support of Intel Cache Allocation Technology. This technology enables the software to restrict cache allocation to a defined subset of cache. The defined subset can overlap with other subsets. (BZ#1288964)

qat updated to the latest upstream version

The qat driver has been updated to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version.

Notable bug fixes and enhancements:

- Added support for the Diffie-Hellman (DH) software
- Added support for Elliptic Curve Diffie-Hellman (ECDH) software
- Added support for Error-correcting Code (ECC) software for curve P-192 and P-256 (BZ#1382849)

Addition of intel-cmt-cat package

The pqos utility provided in this package enables administrators to monitor and manipulate L3 cache to improve utility and performance.

The tool bypasses the kernel API and operates on the hardware directly, this requires that CPU pinning is in use with the target process before use. (BZ#1315489)

i40e now supports trusted and untrusted VFs

This update adds support for both trusted and untrusted virtual functions into the i40e NIC driver. (BZ#1384456)

Kernel support for OVS 802.1ad (QinQ)

This update provides the ability to use two VLAN tags with Open vSwitch (OVS) by enabling the 802.1ad (QinQ) networking standard in kernel. Note that the user-space part of this update is provided by the `openvswitch` package. (BZ#1155732)

Live post-copy migration support for shared memory and hugetlbfs

This update enhances the kernel to enable live post-copy migration to support shared memory and the `hugetlbfs` file system. To benefit from this feature:

- Configure 2MiB huge pages on a host,
- Create a guest VM with 2MiB huge pages,
- Run the guest VM and a stress-test application to test the memory,
- Live-migrate the guest VM with post-copy. (BZ#1373606)

New package: dbxtool

The `dbxtool` package provides a command-line utility and a one-shot `systemd` service for applying UEFI Secure Boot DBX updates. (BZ#1078990)

m1x5 now supports SRIOV-trusted VFs

This update adds support of Single Root I/O Virtualization (SRIOV)-trusted virtual functions (VFs) to the `m1x5` driver. (BZ#1383280)

rwsem performance updates from the 4.9 kernel backported

With this update, most upstream R/W semaphores (`rwsem`) performance related changes up to the Linux kernel version 4.9 have been backported into the Linux kernel while maintaining kernel Application Binary Interface (KABI).

Notable changes include:

- Writer-optimistic spinning, which reduces locking latency and improves locking performance.
- Lock-less waiter wakeup without holding internal spinlock. (BZ#1416924)

getrandom added to the Linux kernel

This update adds the `getrandom` system call to the Linux kernel. As a result, the user space can now request randomness from the same non-blocking entropy pool used by `/dev/urandom`, and the user space can block until at least 128 bits of entropy has been accumulated in that pool. (BZ#1432218)

A new status line, Umask, has been included in /proc/<PID>/status

Previously, it was not possible to read the process umask without modification. Without this change, a library cannot read the umask safely, especially if the main program is multithreaded. The `proc` filesystem (`procfs`) now exposes the umask in the `/proc/<PID>/status` file. The format is `Umask: 0000`, where `0000` is the octal representation of the umask of the task. (BZ#1391413)

Intel® Omni-Path Architecture (OPA) host software

Intel® Omni-Path Architecture (OPA) host software has been fully supported since Red Hat Enterprise Linux 7.3. 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 how to obtain Intel® Omni-Path Architecture documentation, see <https://access.redhat.com/articles/2039623>. (BZ#1459948)

The XTS-AES key verification now meets the FIPS 140-2 requirements

With this update, while running Red Hat Enterprise Linux in FIPS mode and using kernel XTS-AES key verification, the AES key is forced to be different from the tweak key. This ensures that the FIPS 140-2 IG A.9 requirements are met. Additionally, the XEX-based tweaked-codebook mode with ciphertext stealing (XTS) test vectors now could be marked to be skipped. (BZ#1314179)

m1x5 is now supported on IBM z Systems

The Mellanox m1x5 device driver is now also supported for Linux on IBM z Systems and can be used for Ethernet TCP/IP network. (BZ#1394197)

The perf tool now supports processor cache-line contention detection

The `perf` tool now provides the `c2c` subcommand for Shared Data Cache-to-Cache (C2C) analysis. This enables you to inspect cache-line contention and detect both true sharing and false sharing.

Contention occurs when a processor core on a Symmetric Multi Processing (SMP) system modifies data items on the same cache line that is in use by other processors. All other processors using this cache line must then invalidate their copy and request an updated one, which can lead to degraded performance.

The new `c2c` subcommand provides detailed information about the cache lines where contention has been detected, the processes reading and writing the data, the instructions causing the contention, and the Non-Uniform Memory Access (NUMA) nodes involved. (BZ#1391243)

SCSI-MQ support in the lpfc driver

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

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#1382101)

CHAPTER 13. REAL-TIME KERNEL

About Red Hat Enterprise Linux for Real Time Kernel

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

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

kernel-rt rebased

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#1391779)

CHAPTER 14. NETWORKING

NetworkManager rebased to version 1.8

The NetworkManager package has been upgraded to upstream version 1.8, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Support for additional route options has been added.
- Managed state of device until reboot has been persisted.
- Devices that are externally managed are now correctly handled.
- Networked reliability on multihomed hosts has been enhanced.
- Hostname management is now more flexibly configured.
- Support for changing and enforcing 802-3 link properties has been added. (BZ#[1414103](#))

NetworkManager now supports additional features for routes

With this update, NetworkManager can set some advanced options: `source_address` (src, IPv4 only), `from`, `type_of_service` (tos), `window`, `maximum_transmission_unit` (mtu), `congestion_window` (cwnd), `initial_congestion_window` (initcwnd), and `initial_receiver_window` (initrwnd) for static IPv4 and IPv6 routes of connections. (BZ#[1373698](#))

NetworkManager now better handles devices state

With this update, NetworkManager now maintains the state of devices after the service restart and takes over interfaces which are set into managed mode during restart. In addition, NetworkManager can handle devices which are not explicitly set as unmanaged but controlled manually by the user or another network service. (BZ#[1394579](#))

NetworkManager now supports MACsec (IEEE 802.1AE)

This update adds support for configuring Media Access Control Security (MACsec) encryption into NetworkManager. (BZ#[1337997](#))

NetworkManager now supports changing and enforcing 802-3 link properties

Previously, NetworkManager only exposed 802-3 link properties: `802-3-ethernet.speed`, `802-3-ethernet.duplex`, and `802-3-ethernet.auto-negotiate`. With this update, it is possible to change and enforce them. You can either do this automatically using `auto-negotiate=yes`, or manually using `auto-negotiate=no`, `speed=<Mbit/s>`, `duplex=[half, full]`.

Note that if `auto-negotiate=no` and either `speed` or `duplex` are not set, then the link negotiation is skipped and the `auto-negotiate=no`, `speed=0`, `duplex=NULL` default values are preserved.

Note also that the `auto-negotiate` default value has been changed from `yes` to `no` to preserve backward compatibility. Previously, the property was ignored, but now an `auto-negotiate` value of `yes` can enforce link negotiation. Setting it to `no` with `speed` and/or `duplex` unset means that link negotiation is ignored. (BZ#[1353612](#))

NetworkManager now supports ordering bond slaves based on device names

Previously, the existing order of activation for slave connections could cause problems determining the MAC address of the master interface. This update adds more predictable ordering based on device names. You can enable the new ordering using the `slaves-order=name` setting in NetworkManager

configuration.

Note that the new ordering is disabled by default and must be explicitly enabled. (BZ#[1420708](#))

NetworkManager now supports VFs for SR-IOV devices

With this update, the **NetworkManager** system service supports creating virtual functions (VFs) for Single Root I/O Virtualization (SR-IOV) PCI devices. The number of VFs can be specified using the **sriov-num-vfs** option in the device section of the **NetworkManager** configuration file. After VFs are created, **NetworkManager** can activate connection profiles on them.

Note that some properties of a VF interface, such as the Maximum Transmission Unit (MTU), can only be set to values compatible with those that are set on the physical interface. (BZ#[1398934](#))

Kernel GRE rebased to version 4.8

Kernel Generic Routing Encapsulation (GRE) tunneling has been updated to upstream version 4.8, which provides a number of bug fixes and enhancements over the previous version. The most notable changes include:

- Code merge for transmit and receive paths for IPv4 GRE and IPv6 GRE
- Enhancements that allow link layer address changes without bringing the **gre** (IPv4 GRE) or **ip6gre** (IPv6 GRE) device down
- Support for various offloads such as **checksum**, **scatter-gather**, **highdma**, **gso**, or **gro**, for IPv6 GRE traffic
- Automatic kernel module loading when adding **ip6gretap** devices
- Miscellaneous tunneling fixes (such as error handling, MTU calculation, path MTU discovery) up to Linux kernel version 4.8 that affect GRE tunnels (BZ#[1369158](#))

dnsmasq rebased to version 2.76

The **dnsmasq** packages have been upgraded to version 2.76, which provides a number of bug fixes and enhancements. Notable changes include the following:

- The **dhcp_release6** utility is now supported.
- The **ra-param** option has been added.
- Support for the RFC-4242 **information-refresh-time** options in the reply to the DHCPv6 information request has been added.
- The **ra-advrouter** mode for RFC-3775-compliant mobile IPv6 support has been added.
- The **script-arp** script has been added and two new functions for the **dhcp-script** script have been included.
- It is now possible to use random addresses for DHCPv6 temporary address allocations, instead of algorithmically determined stable addresses.
- New optional DNS Security Extensions (DNSSEC) support has been disabled.
- **dnsmasq** can change the default values of IPv6 Router Advertisement. As a result, the **ra-param** option is used to change the default priorities and time intervals of routes advertised by **dnsmasq**. See the **dnsmasq(1)** man page for more information. (BZ#[1375527](#), BZ#[1398337](#))

BIND changes the way it handles URI resource records, impacting also URI backward compatibility

With this update, the BIND suite no longer adds an additional length byte to a value field when using a URI resource record. This also means that BIND in Red Hat Enterprise Linux (RHEL) 7.4 communicates only in the format described in RFC 7553: <https://tools.ietf.org/html/rfc7553>.

Note that this update makes new URI records incompatible with records created using BIND in previous versions of RHEL. Namely, BIND in RHEL 7.4 cannot:

- Understand URI records provided by previous versions of BIND in RHEL.
- Serve URI records to clients using previous versions of BIND in RHEL.

However, BIND in RHEL 7.4 still can:

- Cache and receive records from both earlier and future versions of BIND in RHEL.
- Serve records in the old URI format encoded as Unknown DNS Resource Record. See RFC 3597 for details: <https://tools.ietf.org/html/rfc3597>.

After this update, you do not need to make any change to the DNS zone files. (BZ#[1388534](#))

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

An example of the **DHCP** client hook for Dynamic DNS (DDNS) for Microsoft Azure cloud has been added to the **dhclient** package. The administrator can now easily enable this hook, and register Red Hat Enterprise Linux clients with a **DDNS** server. (BZ#1374119)

dhcp_release6 now releases IPv6 addresses

With this update, the **dhcp_release6** utility can release Dynamic Host Configuration Protocol version 6 (DHCPv6) leases for IPv6 addresses on the local **dnsmasq** server. See the **dhcp_release6(1)** man page for more information about the **dhcp_release6** command. (BZ#[1375569](#))

Sendmail now supports ECDHE

This update adds the Elliptic Curve Diffie-Hellman Ephemeral Keys (ECDHE) support to Red Hat Enterprise Linux 7 **Sendmail**. ECDHE is a variant of the Diffie-Hellman protocol that uses elliptic curve cryptography. It is an anonymous key agreement protocol that allows two parties to establish a shared secret over an insecure channel. (BZ#[1124827](#))

telnet now supports the -6 option

With this update, the **telnet** utility supports the **-6** option to test IPv6 connections. (BZ# [1367415](#))

Adjustable TTL limit for caching negative DNS responses in Unbound

This update adds the **cache-max-negative-ttl** configuration option for the **Unbound** service, which enables adjustment of the maximum TTL specifically for caching negative DNS responses. Previously, this limit was determined by the domain SOA record, or it was automatically the same as the maximum TTL limit for caching all DNS responses, if configured.

Note that if **Unbound** is determining the TTL for DNS response caching, the value set for the **cache-min-ttl** option has precedence over the value specified by **cache-max-negative-ttl**. (BZ#1382383)

The scalability of UDP sockets has been improved

This update improves UDP forward memory accounting and reduces the lock contention of UDP sockets. As a result, the overall ingress throughput of UDP sockets receiving traffic from multiple peers is considerably increased without any outward functional changes. (BZ#1388467)

IP now supports `IP_BIND_ADDRESS_NO_PORT` in the kernel

This update adds the `IP_BIND_ADDRESS_NO_PORT` socket option to the kernel. This allows the kernel to skip L4 tuple reservation when a `bind()` request is used to a port number of `0`. As a result, many simultaneous connections to different destination hosts can be maintained. (BZ#1374498)

IPVS Source Hash scheduling now supports L4 hashing and SH fallback

With this update, the IP Virtual Server (IPVS) Source Hash scheduling algorithm includes:

- L4 hashing
- SH fallback of requests to the next active server in case the destination server has a weight of `0`, which indicates that the destination server is inactive.

As a result, it is now possible to balance the load of requests from one source IP address based on port numbers. Requests to inactive servers no longer time out. (BZ#1365002)

iproute now supports changing bridge port options

With this update, changing bridge port options such as `state`, `priority`, and `cost` have been added to the `iproute` package. As a result, `iproute` can be used as an alternative to the `bridge-utils` package. (BZ#1373971)

New options of Sockets API Extensions for SCTP (RFC 6458) implemented

This update implements options `SCTP_SNDINFO`, `SCTP_NXTINFO`, `SCTP_NXTINFO` and `SCTP_DEFAULT_SNDINFO` to the Sockets API Extensions for the Stream Control Transmission Protocol (RFC 6458).

These new options replace the options `SCTP_SNDRCV`, `SCTP_EXTRCV` and `SCTP_DEFAULT_SEND_PARAM`, which are now deprecated. See also the deprecated functionality section. (BZ#1339791)

ss now supports SCTP sockets list

Previously, the `netstat` utility provided a list of Stream Control Transmission Protocol (SCTP) sockets. With this update, the `ss` utility is able to display the same list. (BZ# [1063934](#))

wpa_supplicant rebased to version 2.6

The `wpa_supplicant` packages have been upgraded to upstream version 2.6, which provides a number of bug fixes and enhancements. Notably, the `wpa_supplicant` utility now supports the Media Access Control Security (MACsec) encryption 802.1AE, which enables MACsec to be used in configuration by default. (BZ#[1404793](#), BZ#1338005)

Linux kernel now contains the switchdev infrastructure and `mlxsw`

This update backports the following functionality into the Linux kernel:

- The Ethernet switch device driver model - the `switchdev` infrastructure; as a result, switch devices can now offload forwarding data plane from the kernel
- The `mlxsw` driver

Switch hardware supported by `mlxsw`:

- Mellanox SwitchX-2 (slow path only)
- Mellanox SwitchIB and SwitchIB-2
- Mellanox Spectrum

Features supported by `m1xsw`:

- Per port jumbo frames, speed setting, state setting, statistics
- Port splitting together with splitter cables
- Port mirroring
- QoS: 802.1p, Data Center Bridging (DCB)
- Access Control Lists (ACLs) using TC flower offloading have been introduced as a Technology Preview

Layer 2 features:

- VLANs
- Spanning Tree Protocol (STP)
- Link Aggregation (LAG) using team or bonding offloading
- Link Layer Discovery Protocol (LLDP)

Layer 3 features:

- Unicast routing

To configure all these features, use standard tools provided by the `iproute` package that has been updated as well. (BZ#[1297841](#), BZ#1275772, BZ#1414400, BZ#1434587, BZ#1434591)

The Linux bridge code rebased to version 4.9

The Linux bridge code has been upgraded to upstream version 4.9, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Support for 802.1ad VLAN filtering and Tx VLAN acceleration
- Support for 802.11 Proxy Address Resolution Protocol (ARP)
- Support for switching offloading by using `switchdev`
- VLAN support for user `mdb` entries
- Support for extended attributes in `mdb` entries
- Support for temporary port router
- Support for per-VLAN statistics
- Support for Internet Group Management Protocol/Multicast Listener Discovery (IGMP/MLD) statistics
- All configuration settings supported by using `sysfs` are now supported by `netlink` as well
- Added per-port flag to control the unknown multicast flood (BZ#1352289)

`bind-dyndb-ldap` rebased to version 11.1

The `bind-dyndb-ldap` package has been upgraded to upstream version 11.1, which provides a number of bug fixes and enhancements over the previous version.

Notably, the `/etc/named.conf` file now uses the new DynDB API. Updating the `bind-dyndb-ldap` package automatically converts the file to the new API style. (BZ#[1393889](#))

DynDB API from the upstream version 9.11.0 of BIND added to Red Hat Enterprise Linux

This update backports the API for the `dyndb` system plug-in, which was introduced in the `bind` package version 9.11.0 in upstream. As a result, the `bind-dyndb-ldap` plug-in in Red Hat Enterprise Linux now uses the new API. The downstream feature `dynamic_db`, which was used in previous releases of Red Hat Enterprise Linux, is no longer supported.

Because the upstream `dyndb` uses a different configuration syntax than the downstream `dynamic_db`, the syntax also changes with this update. However, you do not need to make any manual configuration changes. (BZ#[1393886](#))

tboot rebased to version 1.9.5

The `tboot` packages have been upgraded to upstream version 1.9.5, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- This update adds the 2nd generation of the Link Control Protocol (LCP) creation utility for Trusted Platform Module (TPM) 2.0, as well as a user guide for the updated LCP creation utility.
- A workaround has been implemented to ensure the correct behavior of Intel Platform Trust Technology (PTT) and the Linux PTT driver.
- New fields have been added in the Linux kernel header struct declaration, in order to accommodate for new capabilities of the Linux kernel. (BZ#[1384210](#))

Packages related to rdma consolidated by rebase into rdma-core version 13

The packages related to the `rdma` package have been upgraded and consolidated into a single source package, `rdma-core` version 13. The packages are:

- `rdma`
- `iwpmc`
- `libibverbs`
- `librdmacm`
- `ibacm`
- `libibumad`
- `libocrdma`
- `libmlx4`
- `libmlx5`
- `libhfi1verbs`
- `libi40iw`

- `srp_daemon` (formerly `srptools`)
- `libmthca`
- `libcxgb3`
- `libcxgb4`
- `libnes`
- `libipathverbs`
- `librxe`
- `rdma-ndd`

The following, previously not included, packages have been added as part of the new package `rdma-core`:

- `libqedr`
- `libhns`
- `libvmw_pvrDMA`

All `ibverbs` hardware-specific provider libraries are now bundled in the `libibverbs` sub-package, streamlining installation and preventing possible versioning mismatches. (BZ#1404035)

OVN IP address management support added for static MAC addresses

This update adds support for dynamic IP address assignment with user-specified static MAC addresses. As a result, Open Virtual Network (OVN) users can now create configurations with dynamic IP that are associated with static MAC addresses. (BZ#1368043)

Enhanced networked reliability on multihomed hosts

On interfaces with a route that is already present on another interface, the `NetworkManager` utility now automatically switches the reverse path filtering method from `Strict` to `Loose`. This enhances network reliability on multihomed host machines. (BZ#1394344)

Offloading of GENEVE, VXLAN, and GRE tunnels is now supported

With this update, the infrastructure to support offloading of GENEVE, VXLAN, and GRE tunnels has been added. In addition, various bugs have been fixed in the GENEVE tunnel implementation. (BZ#1326309)

LCO for tunnel traffic is now supported

With this update, the `Local Checksum Offloading` (LCO) technique has been added to enable certain network cards to utilize checksum offloading for tunnel traffic. This enhancement improves the performance of VXLAN, GRE, and other tunnels. (BZ#1326318)

Improved tunnel performance on NICs

With this update, tunnel performance on some Network Interface Cards (NICs) that do not support tunnel offloads by default has been enhanced. As a result, users can now take advantage of existing hardware offloads on these NICs. (BZ#1326353)

NPT is now supported in the kernel

With this update, the **IPv6-to-IPv6 Network Prefix Translation** (NPTv6) function defined in RFC 6296 has been added in the Netfilter framework. As a result, it is now possible to enable **NPT** for stateless translation between IPv6 prefixes. (BZ#1432897)

DNS configuration is now supported through the D-Bus API

Previously, external applications could not easily retrieve the **DNS** parameters used by **NetworkManager**. With this update, DNS configuration has been supported through the D-Bus API. As a result, all DNS-related information, including name servers and domains, is available to client applications through the D-Bus API of **NetworkManager**. An example of such application is the **nmcli** tool, which can now display DNS configuration. (BZ#1404594)

PPP support is now moved into a separate package

With this update, the Point-to-Point Protocol (PPP) support is moved into a separate, optional **NetworkManager-ppp** package. As a result, the dependency chain of **NetworkManager** is smaller and it is possible to limit the number of installed packages.

Note that to configure PPP settings, you must make sure that the **NetworkManager-ppp** package is installed. (BZ#1404598)

The tc utility now supports flower

The **tc** utility has been enhanced to use the kernel **flower** traffic control classifier. With this update, a user can add, modify, or delete **flower** classifier rules from an interface. (BZ# 1422629)

Fix to the CRC32c value computation in SCTP forwarding path

Previously, the kernel incorrectly computed the **CRC32c** value of Stream Control Transmission Protocol (SCTP) packets with offloaded checksum when the kernel forwarded them to an interface that did not support offloading. This update fixes the computation of **CRC32c** in the forwarding path. As a result, SCTP packets are now correctly transmitted in the described situation. (BZ#1072503)

New packages: iperf3

This update adds the **iperf3** packages version 3.1.7 to Red Hat Enterprise Linux 7. The **iperf3** utility enables active measuring of the maximum achievable bandwidth on IP networks. (BZ#913329)

Installation of OVN now supports easily-configurable firewalld rules

This feature adds **firewalld** configuration rules for Open Virtual Network (OVN) to the **openvswitch** packages. As a result, the user can install easier OVN with **firewalld** enabled, instead of needing to create **firewalld** configuration manually. (BZ# 1390938)

netlink now supports bridge master attributes

With this update, whenever bridge attributes are changed, a notification is sent out to listeners. This includes changes triggered by **sysfs**, **rtnl**, **ioctl**, or user applications, such as **NetworkManager**. (BZ#950243)

CHAPTER 15. SECURITY

New packages: tang, clevis, jose, luksmeta

Network Bound Disk Encryption (NBDE) allows the user to encrypt root volumes of the hard drives on physical and virtual machines without requiring to manually enter password when systems are rebooted.

- Tang is a server for binding data to network presence. It includes a daemon which provides cryptographic operations for binding to a remote service. The tang package provides the server side of the NBDE project.
- Clevis is a pluggable framework for automated decryption. It can be used to provide automated decryption of data or even automated unlocking of LUKS volumes. The clevis package provides the client side of the NBDE project.
- José is a C-language implementation of the Javascript Object Signing and Encryption standards. The jose package is a dependency of the clevis and tang packages.
- LUKSMeta is a simple library for storing metadata in the LUKSv1 header. The luksmeta package is a dependency of the clevis and tang packages.

Note that the tang-nagios and clevis-udisk2 subpackages are available only as a Technology Preview. (BZ#1300697, BZ#1300696, BZ#1399228, BZ#1399229)

New package: usbguard

The USBGuard software framework provides system protection against intrusive USB devices by implementing basic whitelisting and blacklisting capabilities based on device attributes. To enforce a user-defined policy, USBGuard uses the Linux kernel USB device authorization feature. The USBGuard framework provides the following components:

- The daemon component with an inter-process communication (IPC) interface for dynamic interaction and policy enforcement
- The command-line interface to interact with a running USBGuard instance
- The rule language for writing USB device authorization policies
- The C++ API for interacting with the daemon component implemented in a shared library (BZ#1395615)

openssh rebased to version 7.4

The openssh package has been updated to upstream version 7.4, which provides a number of enhancements, new features, and bug fixes, including:

- Added support for the resumption of interrupted uploads in SFTP.
- Added the extended log format for the authentication failure messages.
- Added a new fingerprint type that uses the SHA-256 algorithm.
- Added support for using PKCS#11 devices with external PIN entry devices.
- Removed support for the SSH-1 protocol from the OpenSSH server.
- Removed support for the legacy v00 cert format.

- Added the **PubkeyAcceptedKeyTypes** and **HostKeyAlgorithms** configuration options for the **ssh** utility and the **sshd** daemon to allow disabling key types selectively.
- Added the **AddKeysToAgent** option for the **OpenSSH** client.
- Added the **ProxyJump** **ssh** option and the corresponding **-J** command-line flag.
- Added support for key exchange methods for the Diffie-Hellman 2K, 4K, and 8K groups.
- Added the **Include** directive for the **ssh_config** file.
- Removed support for the **UseLogin** option.
- Removed support for the pre-authentication compression in the server.
- The **seccomp** filter is now used for the pre-authentication process. (BZ#1341754)

audit rebased to version 2.7.6

The audit packages have been updated to upstream version 2.7.6, which provides a number of enhancements, new features, and bug fixes, including:

- The **auditd** service now automatically adjusts logging directory permissions when it starts up. This helps keep directory permissions correct after performing a package upgrade.
- The **ausearch** utility has a new **--format** output option. The **--format text** option presents an event as an English sentence describing what is happening. The **--format csv** option normalizes logs into a subject, object, action, results, and how it occurred in addition to some metadata fields which is output in the Comma Separated Value (CSV) format. This is suitable for pushing event information into a database, spreadsheet, or other analytic programs to view, chart, or analyze audit events.
- The **auditctl** utility can now reset the lost event counter in the kernel through the **--reset-lost** command-line option. This makes checking for lost events easier since you can reset the value to zero daily.
- **ausearch** and **aureport** now have a **boot** option for the **--start** command-line option to find events since the system booted.
- **ausearch** and **aureport** provide a new **--escape** command-line option to better control what kind of escaping is done to audit fields. It currently supports **raw**, **tty**, **shell**, and **shell_quote** escaping.
- **auditctl** no longer allows rules with the entry filter. This filter has not been supported since Red Hat Enterprise Linux 5. Prior to this release, on Red Hat Enterprise Linux 6 and 7, **auditctl** moved any entry rule to the exit filter and displayed a warning that the entry filter is deprecated. (BZ#1381601)

opensc rebased to version 0.16.0

The **OpenSC** set of libraries and utilities provides support for working with smart cards. **OpenSC** focuses on cards that support cryptographic operations and enables their use for authentication, mail encryption, or digital signatures.

Notable enhancements in Red Hat Enterprise Linux 7.4 include:

- **OpenSC** adds support for Common Access Card (CAC) cards.

- **OpenSC** implements the **PKCS#11** API and now provides also the **CoolKey** applet functionality. The **opensc** packages replace the **coolkey** packages.

Note that the **coolkey** packages will remain supported for the lifetime of Red Hat Enterprise Linux 7, but new hardware enablement will be provided through the **opensc** packages. (BZ#[1081088](#), BZ#[1373164](#))

openssl rebased to version 1.0.2k

The **openssl** package has been updated to upstream version 1.0.2k, which provides a number of enhancements, new features, and bug fixes, including:

- Added support for the Datagram Transport Layer Security TLS (DTLS) protocol version 1.2.
- Added support for the automatic elliptic curve selection for the ECDHE key exchange in TLS.
- Added support for the Application-Layer Protocol Negotiation (ALPN).
- Added Cryptographic Message Syntax (CMS) support for the following schemes: RSA-PSS, RSA-OAEP, ECDH, and X9.42 DH.

Note that this version is compatible with the API and ABI in the **OpenSSL** library version in previous releases of Red Hat Enterprise Linux 7. (BZ#[1276310](#))

openssl-ibmca rebased to version 1.3.0

The **openssl-ibmca** package has been updated to upstream version 1.3.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- Added support for SHA-512.
- Cryptographic methods are dynamically loaded when the **ibmca** engine starts. This enables **ibmca** to direct cryptographic methods if they are supported in hardware through the **libica** library.
- Fixed a bug in block-size handling with stream cipher modes. (BZ#[1274385](#))

OpenSCAP 1.2 is NIST-certified

OpenSCAP 1.2, the Security Content Automation Protocol (SCAP) scanner, has been certified by the National Institute of Standards and Technology (NIST) as a U. S. government-evaluated configuration and vulnerability scanner for Red Hat Enterprise Linux 6 and 7. **OpenSCAP** analyzes and evaluates security automation content correctly and it provides the functionality and documentation required by NIST to run in sensitive, security-conscious environments. Additionally, **OpenSCAP** is the first NIST-certified configuration scanner for evaluating Linux containers. Use cases include evaluating the configuration of Red Hat Enterprise Linux 7 hosts for PCI and DoD Security Technical Implementation Guide (STIG) compliance, as well as performing known vulnerability scans using Red Hat Common Vulnerabilities and Exposures (CVE) data. (BZ#[1363826](#))

libreswan rebased to version 3.20

The **libreswan** packages have been upgraded to upstream version 3.20, which provides a number of bug fixes and enhancements over the previous version. Notable enhancements include:

- Added support for Opportunistic IPsec (Mesh Encryption), which enables IPsec deployments that cover a large number of hosts using a single simple configuration on all hosts.
- FIPS further tightened.
- Added support for routed-based VPN using Virtual Tunnel Interface (VTI).

- Improved support for non-root configurations.
- Improved Online Certificate Status Protocol (OCSP) and Certificate Revocation Lists (CRL) support.
- Added new **whack** command options: **--fipsstatus**, **--fetchcrls**, **--globalstatus**, and **--shuntstatus**.
- Added support for the NAT Opportunistic Encryption (OE) Client Address Translation: **leftcat=yes**.
- Added support for the Traffic Flow Confidentiality mechanism: **tfc=**.
- Updated cipher preferences as per RFC 4307bis and RFC 7321bis.
- Added support for Extended Sequence Numbers (ESN): **esn=yes**.
- Added support for disabling and increasing the replay window: **replay-window=**. (BZ#[1399883](#))

Audit now supports filtering based on session ID

With this update, the Linux Audit system supports user rules to filter audit messages based on the **sessionid** value. (BZ#1382504)

libseccomp now supports IBM Power architectures

With this update, the **libseccomp** library supports the IBM Power, 64-bit IBM Power, and 64-bit little-endian IBM Power architectures, which enables the GNOME rebase. (BZ#[1425007](#))

AUDIT_KERN_MODULE now records module loading

The **AUDIT_KERN_MODULE** auxiliary record has been added to **AUDIT_SYSCALL** records for the **init_module()**, **fini_module()**, and **delete_module()** functions. This information is stored in the **audit_context** structure. (BZ#1382500)

OpenSSH now uses SHA-2 for public key signatures

Previously, **OpenSSH** used the SHA-1 hash algorithm for public key signatures using RSA and DSA keys. SHA-1 is no longer considered secure, and new SSH protocol extension allows to use SHA-2. With this update, SHA-2 is the default algorithm for public key signatures. SHA-1 is available only for backward compatibility purposes. (BZ#1322911)

firewalld now supports additional IP sets

With this update of the **firewalld** service daemon, support for the following **ipset** types has been added:

- **hash:ip,port**
- **hash:ip,port,ip**
- **hash:ip,port,net**
- **hash:ip,mark**
- **hash:net,net**
- **hash:net,port**

- `hash:net,port,net`
- `hash:net,iface`

The following `ipset` types that provide a combination of sources and destinations at the same time are not supported as sources in `firewalld`. IP sets using these types are created by `firewalld`, but their usage is limited to direct rules:

- `hash:ip,port,ip`
- `hash:ip,port,net`
- `hash:net,net`
- `hash:net,port,net`

The `ipset` packages have been rebased to upstream version 6.29, and the following `ipset` types are now additionally supported:

- `hash:mac`
- `hash:net,port,net`
- `hash:net,net`
- `hash:ip,mark` (BZ#[1419058](#))

firewalld now supports actions on ICMP types in rich rules

With this update, the `firewalld` service daemon allows using Internet Control Message Protocol (ICMP) types in rich rules with the `accept`, `log` and `mark` actions. (BZ#[1409544](#))

firewalld now supports disabled automatic helper assignment

This update of the `firewalld` service daemon introduces support for the disabled automatic helper assignment feature. `firewalld` helpers can be now used without adding additional rules also if automatic helper assignment is turned off. (BZ#1006225)

nss and nss-util now use SHA-256 by default

With this update, the default configuration of the NSS library has been changed to use a stronger hash algorithm when creating digital signatures. With RSA, EC, and 2048-bit (or longer) DSA keys, the SHA-256 algorithm is now used.

Note that also the NSS utilities, such as `certutil`, `crlutil`, and `cmsutil`, now use SHA-256 in their default configurations. (BZ#[1309781](#))

Audit filter exclude rules now contain additional fields

The exclude filter has been enhanced, and it now contains not only the `msgtype` field, but also the `pid`, `uid`, `gid`, `auid`, `sessionID`, and `SELinux` types. (BZ#1382508)

PROCTITLE now provides the full command in Audit events

This update introduces the `PROCTITLE` record addition to Audit events. `PROCTITLE` provides the full command being executed. The `PROCTITLE` value is encoded so it is not able to circumvent the Audit event parser. Note that the `PROCTITLE` value is still not trusted since it is manipulable by the user-space data. (BZ#1299527)

nss-softokn rebased to version 3.28.3

The nss-softokn packages have been upgraded to upstream version 3.28.3, which provides a number of bug fixes and enhancements over the previous version:

- Added support for the ChaCha20-Poly1305 (RFC 7539) algorithm used by TLS (RFC 7905), the Internet Key Exchange Protocol (IKE), and IPsec (RFC 7634).
- For key exchange purposes, added support for the Curve25519/X25519 curve.
- Added support for the Extended Master Secret (RFC 7627) extension. (BZ#[1369055](#))

libica rebased to version 3.0.2

The libica package has been upgraded to upstream version 3.0.2, which provides a number of fixes over the previous version. Notable additions include

- support for Federal Information Processing Standards (FIPS) mode
- support for generating pseudorandom numbers, including enhanced support for Deterministic Random Bit Generator compliant with the updated security specification NIST SP 800-90A. (BZ#1391558)

opencryptoki rebased to version 3.6.2

The opencryptoki packages have been upgraded to upstream version 3.6.2, which provides a number of bug fixes and enhancements over the previous version:

- Added support for **OpenSSL 1.1**
- Replaced deprecated **OpenSSL** interfaces.
- Replaced deprecated **libica** interfaces.
- Improved performance for IBM Crypto Accelerator (ICA).
- Added support for the `rc=8, reasoncode=2028` error message in the `icsf` token. (BZ#1391559)

AUDIT_NETFILTER_PKT events are now normalized

The `AUDIT_NETFILTER_PKT` audit events are now simplified and message fields are now displayed in a consistent manner. (BZ#1382494)

p11tool now supports writing objects by specifying a stored ID

With this update, the `p11tool` GnuTLS PKCS#11 tool supports the new `--id` option to write objects by specifying a stored ID. This allows the written object to be addressable by more applications than `p11tool`. (BZ#[1399232](#))

new package: nss-pem

This update introduces the `nss-pem` package, which previously was part of the `nss` packages, as a separate package. The `nss-pem` package provides the PEM file reader for Network Security Services (NSS) implemented as a PKCS#11 module. (BZ#1316546)

pmrfc3164 replaces pmrfc3164sd in rsyslog

With the update of the `rsyslog` packages, the `pmrfc3164sd` module, which is used for parsing logs in the BSD `syslog` protocol format (RFC 3164), has been replaced by the official `pmrfc3164` module. The official module does not fully cover the `pmrfc3164sd` functionality, and thus it is still available in `rsyslog`. However, it is recommended to use new `pmrfc3164` module wherever possible. The `pmrfc3164sd` module is not supported anymore. (BZ#1431616)

libreswan now supports `right=%opportunisticgroup`

With this update, the `%opportunisticgroup` value for the `right` option in the `conn` part of Libreswan configuration is supported. This allows the opportunistic IPsec with X.509 authentication, which significantly reduces the administrative overhead in large environments. (BZ#1324458)

ca-certificates now meet Mozilla Firefox 52.2 ESR requirements

The Network Security Services (NSS) code and Certificate Authority (CA) list have been updated to meet the recommendations as published with the latest Mozilla Firefox Extended Support Release (ESR). The updated CA list improves compatibility with the certificates that are used in the Internet Public Key Infrastructure (PKI). To avoid certificate validation refusals, Red Hat recommends installing the updated CA list on June 12, 2017. (BZ#1444413)

nss now meets Mozilla Firefox 52.2 ESR requirements for certificates

The Certificate Authority (CA) list have been updated to meet the recommendations as published with the latest Mozilla Firefox Extended Support Release (ESR). The updated CA list improves compatibility with the certificates that are used in the Internet Public Key Infrastructure (PKI). To avoid certificate validation refusals, Red Hat recommends installing the updated CA list on June 12, 2017. (BZ#1444414)

scap-security-guide rebased to version 0.1.33

The scap-security-guide packages have been upgraded to upstream version 0.1.33, which provides a number of bug fixes and enhancements over the previous version. In particular, this new version enhances existing compliance profiles and expands the scope of coverage to include two new configuration baselines:

- Extended support for PCI-DSS v3 Control Baseline
- Extended support for United States Government Commercial Cloud Services (C2S).
- Extended support for Red Hat Corporate Profile for Certified Cloud Providers.
- Added support for the Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) for Red Hat Enterprise Linux 7 profile, aligning to the DISA STIG for Red Hat Enterprise Linux V1R1 profile.
- Added support for the Unclassified Information in Non-federal Information Systems and Organizations (NIST 800-171) profile configures Red Hat Enterprise Linux 7 to the NIST Special Publication 800-53 controls identified for securing Controlled Unclassified Information (CUI).
- Added support for the United States Government Configuration Baseline (USGCB/STIG) profile, developed in partnership with the U. S. National Institute of Standards and Technology (NIST), U. S. Department of Defense, the National Security Agency, and Red Hat.

The USGCB/STIG profile implements configuration requirements from the following documents:

- Committee on National Security Systems Instruction No. 1253 (CNSSI 1253)
- NIST Controlled Unclassified Information (NIST 800-171)
- NIST 800-53 control selections for moderate impact systems (NIST 800-53)
- U. S. Government Configuration Baseline (USGCB)
- NIAP Protection Profile for General Purpose Operating Systems v4.0 (OSPP v4.0)
- DISA Operating System Security Requirements Guide (OS SRG)

Note that several previously-contained profiles have been removed or merged. (BZ#[1410914](#))

CHAPTER 16. SERVERS AND SERVICES

chrony rebased to version 3.1

The chrony package has been upgraded to upstream version 3.1, which provides a number of bug fixes and enhancements over the previous version. Notable enhancements include:

- Added support for software and hardware timestamping for improved accuracy (sub-microsecond accuracy may be possible).
- Improved accuracy with asymmetric network jitter.
- Added support for interleaved mode.
- Added support for configuration and monitoring over Unix domain socket to replace authentication with command key (remote configuration is no longer possible).
- Improved automatic replacement of servers.
- Added orphan mode compatible with the `ntpd` daemon.
- Added response rate limiting for NTP servers.
- Added detailed manual pages, which replace the documentation in the info format. (BZ#1387223)

linuxptp rebased to version 1.8

The linuxptp packages have been upgraded to upstream version 1.8, which provides a number of bug fixes and enhancements over the previous version. Notable enhancements include:

- Added support for hybrid end-to-end (E2E) delay measurements using unicast messages to reduce network traffic in large networks.
- Added support for running a boundary clock (BC) using independent Precision Time Protocol (PTP) hardware clocks.
- Added options to configure Time to Live (TTL) and Differentiated Services Code Point (DSCP) of PTP messages. (BZ#1359311)

tuned rebased to version 2.8.0

The tuned packages have been upgraded to upstream version 2.8.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes include the following:

- CPU partitioning profile has been added.
- Support for cores isolation has been added.
- Support for `initrd` overlays has been added.
- Inheritance has been improved.
- RegExp device matching based on the `udev` device manager has been implemented. (BZ#1388454, BZ#1395855, BZ#1395899, BZ#1408308, BZ#1394965)

logrotate now uses `/var/lib/logrotate/logrotate.status` as the default state file

Previously, the `logrotate` cron job used a modified path to the `logrotate` state file. Consequently, the path used by the cron job did not match the default state file path used by

logrotate itself. To prevent confusion, the default state file path used by **logrotate** has been changed to match the state file path used by **logrotate cron job**. As a result, **logrotate** now uses `/var/lib/logrotate/logrotate.status` as the default state file path in both scenarios. (BZ#1381719)

rsyslog rebased to version 8.24.0

The **rsyslog** utility has been rebased to upstream version 8.24.0, which includes numerous enhancements, new features and bug fixes. Notable improvements include:

- A new core engine has been implemented, offering faster message processing.
- Speed and stability when handling data in the JSON format have been improved.
- The RainerScript configuration format has been selected as default and improved with more options.
- A new `mmexternal` module for manipulation of messages inside **rsyslog** using external applications has been added.
- The `omprog` module has received improvements for better communication with external binaries.
- Modules `imrelp` and `omrelp` now support encrypted transmission using the TLS protocol.
- The `imuxsock` module now supports rule sets for individual sockets, which override the global rule set.
- When the `imuxsock` module is used, rate limiting messages now include PID of the process that causes the rate limiting.
- The TCP server error messages now include the IP address of the remote host.
- The `imjournal` module no longer stops receiving logs after switching to the persistent `journald` configuration.
- Logging to the runtime journal no longer completely stops after a reboot when the machine's clock was set to an earlier time.
- Previously, when the **logrotate** utility with `copytruncate` option was rotating a log file, the `imfile` module might not have read all of the log messages from the file being rotated. As a consequence, these log messages were lost. The `imfile` module has been extended to handle this situation. As a consequence, messages are no longer lost when **logrotate copytruncate** is used on log files.

Customers using custom modules are advised to update their modules for the current **rsyslog** version.

See also the [Deprecated Functionality](#) chapter for information about deprecated **rsyslog** options. (BZ#1313490, BZ#1174345, BZ#1053641, BZ#1196230, BZ#1326216, BZ#1088021, BZ#1419228, BZ#1133687)

New cache configuration options for mod_nss

This update adds new options to control caching of OCSP responses to the `mod_nss` module. The new options allow the user to control:

- Time to wait for OCSP responses

- Size of the OCSP cache
- Minimum and maximum duration for an item's presence in cache, including not caching at all (BZ#1392582)

Database and prefix options have been removed from `nss_pcache`

The `nss_pcache` pin-caching service no longer shares the Network Security Services (NSS) database of the `mod_nss` Apache module because `nss_pcache` does not need access to the tokens. The options for the NSS database and the prefix have been removed and are now handled automatically by `mod_nss`. (BZ#1382102)

New package: `libfastjson`

This update introduces the `libfastjson` library as a replacement of the `json-c` library for `rsyslog`. The limited feature set of `libfastjson` allows for greatly improved performance compared to `json-c`. (BZ#1395145)

`tuned` now supports `initrd` overlays

`tuned` now supports `initrd` overlays, which can extend default (Dracut) `initrd` images. It is supported by the bootloader plugin. The example shows typical usage in the Tuned profile:

```
[bootloader]
initrd_add_dir=${i:PROFILE_DIR}/overlay.img
```

This adds the content of the `overlay.img` directory to the current `initrd` when the profile is activated. (BZ#1414098)

`openwsman` now supports disabling of particular SSL protocols

Previously, there was no way to disable particular SSL protocols with the `openwsman` utility. A new configuration file option for a list of disabled protocols has been added. As a result, it is now possible to disable particular SSL protocols through the `openwsman` configuration file. (BZ#1190689)

`rear` rebased to version 2.0

Updated `rear` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 7. Notable changes include:

- The **Cyclic Redundancy Check (CRC)** feature is now enabled by default on the XFS file systems. Previously, `rear` ignored this change in behavior, and formatted the `/boot` partition with an incompatible UUID flag. This caused the recovery process to fail. With this rebase, `rear` checks for the CRC feature, and properly preserves UUID during recovery.
- Support for the **GRUB** and **GRUB2** boot loaders for IBM Power Systems architecture has been added.
- Linux capabilities are now preserved if the directive **NETFS_RESTORE_CAPABILITIES** is set to the `y` option in the `/usr/share/rear/conf/default.conf` configuration file.
- CIFS credentials are now preserved in rescue image.
- **GRUB_SUPERUSER** and **GRUB_RESCUE_PASSWORD** directives have been dropped to avoid possible unexpected behaviour change of the **GRUB2** bootloader in the currently running system.
- Documentation has been improved.

- Creation of multiple backups have been enabled. (BZ#[1355667](#))

python-tornado rebased to version 4.2.1

The python-tornado package has been upgraded to upstream version 4.2.1, which provides a number of bug fixes and new features over the previous version. Notable changes include:

- A new `tornado.netutil.Resolver` class, which provides an asynchronous interface to DNS resolution
- A new `tornado.tcpclient` module, which creates TCP connections with non-blocking DNS, SSL handshaking, and support for IPv6
- The `IOLoop.instance()` function is now thread-safe
- Logging has been improved; low-level logs are less frequent; **Tornado** uses its own logger instead of the root logger, which enables more detailed configuration
- Multiple reference cycles have been separated within python-tornado, enabling more efficient garbage collection on **CPython**
- Coroutines are now faster and are used extensively within **Tornado**. (BZ#[1158617](#))

CHAPTER 17. STORAGE

Support added in LVM for RAID level takeover

LVM now provides full support for RAID takeover, previously available as a Technology Preview, which allows users to convert a RAID logical volume from one RAID level to another. This release expands the number of RAID takeover combinations. Support for some transitions may require intermediate steps. New RAID types that are added by means of RAID takeover are not supported in older released kernel versions; these RAID types are `raid0`, `raid0_meta`, `raid5_n`, and `raid6_{ls,rs,la,ra,n}_6`. Users creating those RAID types or converting to those RAID types on Red Hat Enterprise Linux 7.4 cannot activate the logical volumes on systems running previous releases. RAID takeover is available only on top-level logical volumes in single machine mode (that is, takeover is not available for cluster volume groups or while the RAID is under a snapshot or part of a thin pool). (BZ#[1366296](#))

LVM now supports RAID reshaping

LVM now provides support for RAID reshaping. While takeover allows users to change from one RAID type to another, reshaping allows users to change properties such as the RAID algorithm, stripe size, region size, or number of images. For example, a user can change a 3-way stripe to a 5-way stripe by adding two additional devices. Reshaping is available only on top-level logical volumes in single machine mode, and only while the logical volume is not in-use (for example, when it is mounted by a file system). (BZ#[1191935](#), BZ#[834579](#), BZ#[1191978](#), BZ#[1392947](#))

Device Mapper linear devices now support DAX

Direct Access (DAX) support has been added to the `dm-linear` and `dm-stripe` targets. Multiple Non-Volatile Dual In-line Memory Module (NVDIMM) devices can now be combined to provide larger persistent memory (PMEM) block devices. (BZ#[1384648](#))

libstoragemgmt rebased to version 1.4.0

The `libstoragemgmt` packages have been upgraded to upstream version 1.4.0, which provides a number of bug fixes and enhancements over the previous version. Notably, the following libraries have been added:

- Query serial number of local disk: `lsm_local_disk_serial_num_get()/lsm.LocalDisk.serial_num_get()`
- Query LED status of local disk: `lsm_local_disk_led_status_get()/lsm.LocalDisk.led_status_get()`
- Query link speed of local disk: `lsm_local_disk_link_speed_get()/lsm.LocalDisk.link_speed_get()`

Notable bug fixes include:

- The `megaraid` plug-in for the Dell PowerEdge RAID Controller (PERC) has been fixed.
- The local disk rotation speed query on the NVM Express (NVMe) disk has been fixed.
- `lsmcli` incorrect error handling on a local disk query has been fixed.
- All gcc compile warnings have been fixed.
- The obsolete usage of the `autoconf AC_OUTPUT` macro has been fixed. (BZ# [1403142](#))

mpt3sas updated to version 15.100.00.00

The `mpt3sas` storage driver has been updated to version 15.100.00.00, which adds support for new devices. Contact your vendor for more details. (BZ#[1306453](#))

The `lpfc_no_hba_reset` module parameter for the `lpfc` driver is now available

With this update, the `lpfc` driver for certain models of Emulex Fibre Channel Host Bus Adapters (HBAs) has been enhanced by adding the `lpfc_no_hba_reset` module parameter. This parameter accepts a list of one or more hexadecimal world-wide port numbers (WWPNs) of HBAs that are not reset during SCSI error handling.

Now, `lpfc` allows you to control which ports on the HBA may be reset during SCSI error handling time. Also, `lpfc` now allows you to set the `eh_deadline` parameter, which represents an upper limit of the SCSI error handling time. (BZ#1366564)

LVM now detects Veritas Dynamic Multi-Pathing systems and no longer accesses the underlying device paths directly

For LVM to work correctly with Veritas Dynamic Multi-Pathing, you must set `obtain_device_list_from_udev` to 0 in the devices section of the configuration file `/etc/lvm/lvm.conf`. These multi-pathed devices are not exposed through the standard `udev` interfaces and so without this setting LVM will be unaware of their existence. (BZ#1346280)

The libnvdimm kernel subsystem now supports PMEM subdivision

Intel's Non-Volatile Dual In-line Memory Module (NVDIMM) label specification has been extended to allow more than one Persistent Memory (PMEM) namespace to be configured per region (interleave set). The kernel shipped with Red Hat Enterprise Linux 7.4 has been modified to support these new configurations.

Without subdivision support, a single region could previously be used in only one mode: `pmem`, `device dax`, or `sector`. With this update, a single region can be subdivided, and each subdivision can be configured independently of the others. (BZ#1383827)

Warning messages when multipathd is not running

Users now get warning messages if they run a `multipath` command that creates or lists multipath devices while `multipathd` is not running.

If `multipathd` is not running, then the devices are not able to restore paths that have failed or react to changes in the device setup. The `multipathd` daemon now prints a warning message if there are multipath devices and `multipathd` is not running. (BZ# [1359510](#))

c library interface added to multipathd to give structured output

Users can now use the `libdmmp` library to get structured information from `multipathd`. Other programs that want to get information from `multipathd` can now get this information without running a command and parsing the results. (BZ#[1430097](#))

New remove retries multipath configuration value

If a multipath device is temporarily in use when multipath tries to remove it, the remove will fail. It is now possible to control the number of times that the `multipath` command will retry removing a multipath device that is busy by setting the `remove_retries` configuration value. The default value is 0, in which case multipath will not retry failed removes. (BZ#[1368211](#))

New multipathd reset multipaths stats commands

Multipath now supports two new `multipathd` commands: `multipathd reset multipaths stats` and `multipathd reset multipath dev stats`. These commands reset the device stats that `multipathd` tracks for all the devices, or the specified device, respectively. This allows users to reset their device stats after they make changes to them. (BZ#[1416569](#))

New disable_changed_wwids mulitpath configuration parameter

Multipath now supports a new `multipath.conf` defaults section parameter, `disable_changed_wwids`. Setting this will make multipathd notice when a path device changes its `wwid` while in use, and will disable access to the path device until its `wwid` returns to its previous value.

When the `wwid` of a `scsi` device changes, it is often a sign that the device has been remapped to a different LUN. If this happens while the `scsi` device is in use, it can lead to data corruption. Setting the `disable_changed_wwids` parameter will warn users when the `scsi` device changes its `wwid`. In many cases `multipathd` will disable access to the path device as soon as it gets unmapped from its original LUN, removing the possibility of corruption. However `multipathd` is not always able to catch the change before the `scsi` device has been remapped, meaning there may still be a window for corruption. Remapping in-use `scsi` devices is not currently supported. (BZ#1169168)

Updated built-in configuration for HPE 3PAR array

The built-in configuration for the 3PAR array now sets `no_path_retry` to 12. (BZ#1279355)

Added built-in configuration for NFINIDAT InfiniBox.* devices

Multipath now autoconfigures NFINIDAT InfiniBox.* devices (BZ#1362409)

device-mapper-multipath now supports the max_sectors_kb configuration parameter

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

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

Using the `max_sectors_kb` `multipath.conf` parameter is an easy way to set these values before a multipath device is created on top of the path devices, and prevent invalid-sized I/O operations from being passed down. (BZ#1394059)

New detect_checker multipath configuration parameter

Some devices, such as the VNX2, can be optionally configured in ALUA mode. In this mode, they need to use a different `path_checker` and `prioritizer` than in their non-ALUA mode. Multipath now supports the `detect_checker` parameter in the `multipath.conf` defaults and devices sections. If this is set, multipath will detect if a device supports ALUA, and if so, it will override the configured `path_checker` and use the TUR checker instead. The `detect_checker` option allows devices with an optional ALUA mode to be correctly autoconfigured, regardless of what mode they are in. (BZ#1372032)

Multipath now has a built-in default configuration for Nimble Storage devices

The multipath default hardware table now includes an entry for Nimble Storage arrays. (BZ#1406226)

LVM supports reducing the size of a RAID logical volume

As of Red Hat Enterprise Linux 7.4, you can use the `lvreduce` or `lvresize` command to reduce the size of a RAID logical volume. (BZ#1394048)

iprutils rebased to version 2.4.14

The `iprutils` packages have been upgraded to upstream version 2.4.14, which provides a number of bug fixes and enhancements over the previous version. Notably:

- Endian swapped `device_id` is now compatible with earlier versions.
- VSET write cache in bare metal mode is now allowed.

- Creating RAIDS on dual adapter setups has been fixed.
- Verifying rebuilds for single adapter configurations is now disabled by default. (BZ#1384382)

mdadm rebased to version 4.0

The mdadm packages have been upgraded to upstream version 4.0, which provides a number of bug fixes and enhancements over the previous version. Notably, this update adds bad block management support for Intel Matrix Storage Manager (IMSM) metadata. The features included in this update are supported on external metadata formats, and Red Hat continues supporting the Intel Rapid Storage Technology enterprise (Intel RSTe) software stack. (BZ#1380017)

LVM extends the size of a thin pool logical volume when a thin pool fills over 50 percent

When a thin pool logical volume fills by more than 50 percent, by default the `dmeventd thin` plugin now calls the `dmeventd thin_command` command with every 5 percent increase. This resizes the thin pool when it has been filled above the configured `thin_pool_autoextend_threshold` in the `activation` section of the configuration file. A user may override this default by configuring an external command and specifying this command as the value of `thin_command` in the `dmeventd` section of the `lvm.conf` file. For information on the `thin` plugin and on configuring external commands to maintain a thin pool, see the `dmeventd(8)` man page.

In previous releases, when a thin pool resize failed, the `dmeventd` plugin would try to unmount unconditionally all thin volumes associated with the thin pool when a compile-time defined threshold of more than 95 percent was reached. The `dmeventd` plugin, by default, no longer unmounts any volumes. Reproducing the previous logic requires configuring an external script. (BZ#1442992)

LVM now supports dm-cache metadata version 2

LVM/DM cache has been significantly improved. It provides support for larger cache sizes, better adaptation to changing workloads, greatly improved startup and shutdown times, and higher performance overall. Version 2 of the dm-cache metadata format is now the default when creating cache logical volumes with LVM. Version 1 will continue to be supported for previously created LVM cache logical volumes. Upgrading to version 2 will require the removal of the old cache layer and the creation of a new cache layer. (BZ#1436748)

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

SCSI T10 DIF/DIX is fully supported in Red Hat Enterprise Linux 7.4, 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

QLOGIC 16G FC HBA:

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

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

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

EMC supports DIF on:

EMULEX 8G FC HBA:

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

EMULEX 16G FC HBA:

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

QLOGIC 16G FC HBA:

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

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

Support for DIF/DIX remains in Technology Preview for other HBAs and storage arrays. (BZ#1457907)

The `dmstats` facility can now track the statistics for files that change

Previously, the `dmstats` facility was able to report statistics for files that did not change in size. It now has the ability to watch files for changes and update its mappings to track file I/O even as the file changes in size (or fills holes that may be in the file). (BZ#1378956)

Support for thin snapshots of cached logical volumes

LVM in Red Hat Enterprise Linux 7.4 allows you to create thin snapshots of cached logical volumes. This feature was not available in earlier releases. These external origin cached logical volumes are converted to a read-only state and thus can be used by different thin pools. (BZ#1189108)

New package: `nvmetcli`

The `nvmetcli` utility enables you to configure Red Hat Enterprise Linux as an NVMeoF target, using the NVMe-over-RDMA fabric type. With `nvmetcli`, you can configure `nvmet` interactively, or use a JSON file to save and restore the configuration. (BZ#1383837)

Device DAX is now available for NVDIMM devices

Device DAX enables users like hypervisors and databases to have raw access to persistent memory without an intervening file system. In particular, Device DAX allows applications to have predictable fault granularities and the ability to flush data to the persistence domain from user space. Starting with Red Hat Enterprise Linux 7.4, Device Dax is available for Non-Volatile Dual In-line Memory Module (NVDIMM) devices. (BZ#1383489)

CHAPTER 18. SYSTEM AND SUBSCRIPTION MANAGEMENT

New `payload_gpgcheck` option added to `yum`

With this update, the new configuration option `payload_gpgcheck` has been added to the `yum` utility. This option enables a GNU Privacy Guard (GPG) signature check on the payload sections of packages, thus enhancing the security and integrity when installing packages. Previously, when `gpgcheck` option was enabled, `yum` only performed a GPG signature check on headers. Consequently, if the payload data were tampered with or corrupted, RPM unpacking error occurred, and the package was left in a partly installed state. This might have put the operating system into an inconsistent and vulnerable state.

You can use the new `payload_gpgcheck` option in conjunction with the `gpgcheck` or `localpkg_gpgcheck` options to prevent this problem. As a result, when `payload_gpgcheck` is enabled, `yum` performs a GPG signature check on the payload and aborts the transaction if it is not verified. Using `payload_gpgcheck` is equivalent to manually running `rpm -K` on downloaded packages. (BZ#[1343690](#))

A no-proxy configuration is available for `virt-who`

With this update, the `virt-who` service can be set to ignore proxy network settings. This enables `virt-who` to work properly on environments that use a proxy connection with one-way communication.

To set up this functionality, add the `NO_PROXY` environment variable to the `/etc/sysconfig/virt-who` file. Alternatively, you can add the `no_proxy` variable to the `[server]` section of the `/etc/rhsm/rhsm.conf` file.

Note that the `NO_PROXY` setting does not work when synchronizing the hypervisor using Red Hat Satellite 5. (BZ#[1299643](#))

`virt-who` respects independent interval settings

With this update, the `virt-who` command reports each interval on all sources that have updates. In addition, if `virt-who` is configured to send updates to more than one destination, for example to an Red Hat Satellite instance and the Red Hat Subscription Management (RHSM), the interval for each is maintained separately. This means that all updates can be sent to each configured destination, regardless of the state of communication with other destinations. (BZ#[1436811](#))

Password options added to `virt-who-password`

With this update, the `-p` and `--password` options have been added to the `virt-who-password` utility. This enables the utility to be used in scripts. (BZ#[1426058](#))

Regular expressions and wildcards can be used in some `virt-who` configuration parameters

With this update, regular expressions and wildcards can be used in the `filter_hosts` and `exclude_hosts` configuration parameters. This enables users of `virt-who` to maintain a list of hosts to report on with much more ease.

By using regular expressions and wildcards to specify which hosts to report on or exclude, the hosts list can be much more concise. (BZ#[1405967](#))

`virt-who` configuration files are easier to manage

The `virt-who` service now only uses configuration files in the `/etc/virt-who.d/` directory that end with the `.conf` extension. This enables easier management of `virt-who` configuration files, for example for testing or backup purposes. (BZ#[1369107](#))

CHAPTER 19. VIRTUALIZATION

ENA drivers for Amazon Web Services

This update adds support for Amazon Elastic Network Adapter (ENA) drivers to the Red Hat Enterprise Linux 7 kernel. ENA significantly enhances networking efficiency of Red Hat Enterprise Linux 7 guest virtual machines for certain instance types of the Amazon Web Services cloud.

For more information about ENA, see <https://aws.amazon.com/blogs/aws/elastic-network-adapter-high-performance-network-interface-for-amazon-ec2>. (BZ#1357491, BZ#1410047)

Synthetic Hyper-V FC adapters are supported by the storvsc driver

This update improves the way the `storvsc` driver handles Fibre Channel (FC) devices on Hyper-V virtualization. Notably, when a new synthetic Fibre Channel (FC) adapter is configured on a Hyper-V hypervisor, a new `hostX` (for example `host1`) file is created in the `/sys/class/fc_host/` and `/sys/class/scsi_host/` directories. This file contains the `port_name` and `host_name` entries determined by the Hyper-V FC Adapter world-wide port number (WWPN) and world-wide node number (WWNN). (BZ#1308632, BZ#1425469)

Parent HBA can be defined by a WWNN/WWPN pair

With this release, a parent host bus adapter (HBA) can be identified by a World Wide Node Name (WWNN) and World Wide Port Name (WWPN), in addition to a `scsi_host#`. When defined by a `scsi_host#`, if hardware is added to the host machine, the `scsi_host#` may change after the host machine reboots. By using a WWNN/WWPN pair, the assignment remains unchanged regardless of hardware changes to the host machine. (BZ#1349696)

libvirt rebased to version 3.2.0

The libvirt packages have been upgraded to upstream version 3.2.0, which provides a number of bug fixes and enhancements over the previous version. Notable changes:

- This update makes it possible to install and uninstall specific `libvirt` storage sub-drivers, which reduces the installation footprint.
- You can now configure the `/etc/nsswitch.conf` file to instruct Name Services Switch (NSS) to automatically resolve names of KVM guests to their network addresses. (BZ#1382640)

KVM now supports MCE

This update adds support for Machine Check Exception (MCE) to the KVM kernel modules, which makes it possible to use the Local MCE (LMCE) feature of Intel Xeon v5 processors in KVM guest virtual machines. LMCE can deliver MCE to a single processor thread instead of broadcasting to all threads, which ensures the machine check does not impact the performance of more vCPUs than needed. As a result, this reduces software load when processing MCE on machines with a large number of processor threads. (BZ#1402102, BZ#1402116)

Added support for rx batching on tun/tap devices

With this release, rx batching for tun/tap devices is now supported. This enables receiving bundled network frames which can improve performance. (BZ#1414627)

libguestfs rebased to version 1.36.3

The libguestfs packages have been upgraded to upstream version 1.36.3, which provides a number of bug fixes and enhancements over the previous version. Notable changes include:

- This update adds the `virt-tail` utility, which can be used to follow (tail) log files within a guest, similar to the `tail -f` command. For details, see the `virt-tail(1)` man page.

- The **virt-v2v** utility supports more operating systems and more input sources. In addition, the conversion of Windows guests has been substantially rewritten and simplified.
- Multiple options have been added for the **virt-customize**, **virt-builder**, and **virt-sysprep** utilities. (BZ#[1359086](#))

Improved virt-v2v installation of QXL drivers

This update reworks the **virt-v2v** implementation of QXL driver installation in Windows guest virtual machines, which ensures that QXL drivers are installed correctly on these guests. (BZ#[1233093](#), BZ#[1255610](#), BZ#[1357427](#), BZ#[1374651](#))

virt-v2v can export disk images to qcow2 format 1.1

With this update, the **virt-v2v** utility exports disk images compatible with qcow2 format version 1.1 when using the **-o rhev** option. In addition, **virt-v2v** adds the **--vdsml-compat=COMPAT** option for the vdsml output mode. This option specifies which version of the qcow2 format **virt-v2v** uses when exporting images with the **-o vdsml** option. (BZ#[1400205](#))

Additional virt tools can work on LUKS whole-disk encrypted guests

This update adds support for working on LUKS whole-disk encrypted guests using the **virt-customize**, **virt-get-kernel**, **virt-sparsify**, and **virt-sysprep** tools. As a result, these tools can provide keys or passphrases for opening LUKS whole-disk encrypted guests. (BZ#[1362649](#))

Tab completion for all libguestfs commands

Bash completion scripts have been added for all **libguestfs** tools. As a result, it is now possible to use Tab completion in bash with every **libguestfs** command. (BZ#[1367738](#))

Resized disks can be written directly to a remote location

With this update, the **virt-resize** utility can write its output to a remote location. This may be useful, for example, in directly writing the resized disk image to a Ceph storage volume. The **virt-resize** output disk can be specified using a URI. Any supported input protocol and format can be used to specify the output. (BZ#[1404182](#))

User namespace is now fully supported

The user namespace feature, previously available as a Technology Preview, is now fully supported. It provides additional security to servers running Linux containers by providing better isolation between the host and the containers. Administrators of a container are no longer able to perform administrative operations on the host, which increases security. (BZ#[1138782](#))

Driver added for devices that connect over a PCI Express bus in guest virtual machine under Hyper-V

In this update, a new driver was added that exposes a root PCI bus when a devices that connects over a PCI Express bus is passed through to a Red Hat Enterprise Linux guest virtual machine running on the Hyper-V hypervisor. The feature is currently supported with Microsoft Windows Server 2016. (BZ#[1302147](#))

CHAPTER 20. 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 21. RED HAT SOFTWARE COLLECTIONS

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

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

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



IMPORTANT

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

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

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

PART II. NOTABLE BUG FIXES

This part describes bugs fixed in Red Hat Enterprise Linux 7.4 that have a significant impact on users.

CHAPTER 22. GENERAL UPDATES

Addition of CtrlAltDelBurstAction for Systemd

The `systemd` response to multiple `CTRL+ALT+DEL` events is now configurable by setting the `CtrlAltDelBurstAction` option in `/etc/systemd/system.conf` (BZ#1353028)

cgred can now resolve rules concerning NSS users and groups

Previously, the `cgred` service was not configured to start up after services providing Name Service Switch (NSS) users and groups. Also, information about skipping invalid rules was shown only in debug mode. Consequently, rules in the `cgrules.conf` file concerning NSS users and groups were sometimes ignored without any log message. With this update, `cgred` is configured to start after the `nss-user-lookup` target and level of log messages about skipping rules is changed to warning, which is also set as a default log level for the `cgred` daemon. As a result, NSS users and groups are now always resolved before starting `cgred`. Also, the warning message is logged in case some rules in `cgrules.conf` are invalid. (BZ#1406927)

CHAPTER 23. AUTHENTICATION AND INTEROPERABILITY

yum no longer reports package conflicts after installing ipa-client

After the user installed the ipa-client package, the yum utility unexpectedly reported package conflicts between the ipa and freeipa packages. These errors occurred after failed transactions or after using the `yum check` command. With this update, yum no longer reports errors about self-conflicting packages because such conflicts are allowed by RPM. As a result, yum no longer displays the described errors after installing ipa-client. (BZ#1370134)

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

The Red Hat Directory Server previously tried to retrieve the key slot from a fixed token name, when FIPS mode was enabled on the security database. However, the token name can change. If the key slot is not found, Directory Server is unable to decode the replication manager's password and replication sessions fail. To fix the problem, the `slapd_pk11_getInternalKeySlot()` function now uses FIPS mode to retrieve the current key slot. As a result, replication sessions using SSL or STARTTLS no longer fail in the described situation. (BZ#1378209)

Certificate System no longer fails to install with a Thales HSM on systems in FIPS mode

After installing with the Certificate System (CS) with a Thales hardware security module (HSM), the SSL protocol did not work correctly if you generated all system keys on the HSM. Consequently, CS failed to install on systems with FIPS mode enabled, requiring you to manually modify the `sslRangeCiphers` parameter in the `server.xml` file. This bug has been fixed, and installation FIPS-enabled systems with Thales HSM works as expected. (BZ#1382066)

The dependency list for pkispawn now correctly includes openssl

Previously, when the openssl package was not installed, using the `pkispawn` utility failed with the following error:

```
Installation failed: [Errno 2] No such file or directory
```

This problem occurred because the openssl package was not included as a runtime dependency of the pki-server package contained within the pki-core package. This bug has been fixed by adding the missing dependency, and pkispawn installations no longer fail due to missing openssl. (BZ#1376488)

Error messages from the PKI Server profile framework are now passed through to the client

Previously, PKI Server did not pass through certain error messages generated by the profile framework for certificate requests to the client. Consequently, the error messages displayed on the web UI or in the output of the `pki` command did not describe why a request failed. The code has been fixed and now passes through error messages. Now users can see the reason why an enrollment failed or was rejected. (BZ#1249400)

Certificate System does not start a Lightweight CA key replication during installation

Previously, Certificate System incorrectly started a Lightweight CA key replication during a two-step installation. As a consequence, the installation failed and an error was displayed. With this update, the two-step installation does not start the Lightweight CA key replication and the installation completes successfully. (BZ#1378275)

PKI Server now correctly compares subject DN's during startup

Due to a bug in the routine that adds a Lightweight CA entry for the primary CA, PKI Server previously

failed to compare subject distinguished names (DN) if it contained attributes using encodings other than `UTF8String`. As a consequence, every time the primary CA started, an additional Lightweight CA entry was added. PKI Server now compares the subject DNs in canonical form. As a result, PKI server no longer adds additional Lightweight CA entries in the mentioned scenario. (BZ#[1378277](#))

KRA installation no longer fails when connecting to an intermediate CA with an incomplete certificate chain

Previously, installing a Key Recovery Authority (KRA) subsystem failed with an `UNKNOWN_ISSUER` error if the KRA attempted to connect to an intermediate CA that had a trusted CA certificate but did not have the root CA certificate. With this update, KRA installation ignores the error and completes successfully. (BZ#[1381084](#))

The `startTime` field in certificate profiles now uses long integer format

Previously, Certificate System stored the value in the `startTime` field of a certificate profile as `integer`. If you entered a larger number, Certificate System interpreted the value as a negative number. Consequently, the certificate authority issued certificates that contained a start date located in the past. With this update, the input format of the `startTime` field has been changed to a long integer. As a result, the issued certificates now have a correct start date. (BZ#[1385208](#))

Subordinate CA installation no longer fails due with a `PKCS#11 token is not logged in error`

Previously, subordinate Certificate Authority (sub-CA) installation failed due to a bug in the Network Security Services (NSS) library, which generated the `SEC_ERROR_TOKEN_NOT_LOGGED_IN` error. This update adds a workaround to the installer which allows the installation to proceed. If the error is still displayed, it can now be ignored. (BZ#[1395817](#))

The `pkispawn` script now correctly sets the ECC key sizes

Previously, when a user ran the `pkispawn` script with an Elliptic Curve Cryptography (ECC) key size parameter set to a different value than the default, which is `nistp256`, the setting was ignored. Consequently, the created PKI Server instance issued system certificates, which incorrectly used the default ECC key curve. With this update, PKI Server uses the value set in the `pkispawn` configuration for the ECC key curve name. As a result, the PKI Server instance now uses the ECC key size set when setting up the instance. (BZ#[1397200](#))

CA clone installation in FIPS mode no longer fails

Previously, installing a CA clone or a Key Recovery Authority (KRA) failed in FIPS mode due to an inconsistency in handling internal NSS token names. With this update, the code that handles the token name has been consolidated to ensure that all token names are handled consistently. T allows the KRA and CA clone installation to complete properly in FIPS mode. (BZ#[1411428](#))

PKI Server no longer fails to start when an `entryUSN` attribute contains a value larger than 32-bit

Previously, the `*LDAP Profile Monitor` and the `Lightweight CA Monitor` parsed values in `entryUSN` attributes as a 32-bit integer. As a consequence, when the attribute contained a value larger than that, a `NumberFormatException` error was logged and the server failed to start. The problem has been fixed, and the server no longer fails to start in the mentioned scenario. (BZ#[1412681](#))

Tomcat now works with IPv6 by default

The IPv4-specific `127.0.0.1` loopback address was previously used in the default server configuration file as the default `AJP` host name. This caused connections to fail on servers which run in IPv6-only environments. With this update, the default value is changed to `localhost`, which works with both IPv4 and IPv6 protocols. Additionally, an upgrade script is available to automatically change the `AJP` host name on existing server instances. (BZ# [1413136](#))

pkispawn no longer generates invalid NSS database passwords

Prior to this update, `pkispawn` generated a random password for the NSS database which in some cases contained a backslash (\) character. This caused problems when NSS established SSL connections, which in turn caused the installation to fail with a `ACCESS_SESSION_ESTABLISH_FAILURE` error.

This update ensures that the randomly generated password can not contain the backslash character and a connection can always be established, allowing the installation to finish successfully. (BZ#[1447762](#))

Certificate retrieval no longer fails when adding a user certificate with the `--serial` option

Using the `pki user-cert-add` command with the `--serial` parameter previously used an improperly set up SSL connection to the certificate authority (CA), causing certificate retrieval to fail. With this update, the command uses a properly configured SSL connection to the CA, and the operation now completes successfully. (BZ#[1246635](#))

CA web interface no longer shows a blank certificate request page if there is only one entry

Previously, when the certificate request page in the CA web user interface only contained one entry, it displayed an empty page instead of showing the single entry. This update fixes the web user interface, and the certificate request page now correctly shows entries in all circumstances. (BZ#[1372052](#))

Installing PKI Server in a container environment no longer displays a warning

Previously, when installing the `pki-server` RPM package in a container environment, the `systemd` daemon was reloaded. As a consequence, a warning was displayed. A patch has been applied to reload the daemon only during an RPM upgrade. As a result, the warning is no longer displayed in the mentioned scenario. (BZ#[1282504](#))

Re-enrolling a token using a G&D smart card no longer fails

Previously, when re-enrolling a token using a Giesecke & Devrient (G&D) smart card, the enrollment of the token could fail in certain situations. The problem has been fixed, and as a result, re-enrolling a token works as expected. (BZ#[1404881](#))

PKI Server provides more detailed information about certificate validation errors on startup

Previously, PKI Server did not provide sufficient information if a certificate validation error occurred when the server was started. Consequently, troubleshooting the problem was difficult. PKI Server now uses the new Java security services (JSS) API which provides more detailed information about the cause of the error in the mentioned scenario. (BZ#[1330800](#))

PKI Server no longer fails to re-initialize the `LDAPProfileSubsystem` profile

Due to a race condition during re-initializing the `LDAPProfileSubsystem` profile, PKI Server previously could incorrectly report that the requested profile does not exist. Consequently, requests to use the profile could fail. The problem has been fixed, and requests to use the profile no longer fail. (BZ#[1376226](#))

Extracting private keys generated on an HSM no longer fails

Previously, when generating asymmetric keys on a Lunasa or Thales hardware security module (HSM) using the new Asymmetric Key Generation REST service on the key recovery agent (KRA), PKI Server set incorrect flags. As a consequence, users were unable to retrieve the generated private keys. The code has been updated to set the correct flags for keys generated on these HSMs. As a result, users can now retrieve private keys in the mentioned scenario. (BZ#[1386303](#))

pkispawn no longer generates passwords consisting only of digits

Previously, `pkispawn` could generate a random password for NSS database consisting only digits. Such passwords are not FIPS-compliant. With this update, the installer has been modified to generate FIPS-compliant random passwords which consist of a mix of digits, lowercase letters, uppercase letters, and certain punctuation marks. (BZ#1400149)

CA certificates are now imported with correct trust flags

Previously, the `pki client-cert-import` command imported CA certificates with `CT, c`, trust flags, which was insufficient and inconsistent with other PKI tools. With this update, the command has been fixed and now sets the trust flags for CA certificates to `CT, C, C`. (BZ#1458429)

Generating a symmetric key no longer fails when using the `--usage verify` option

The `pki` utility checks a list of valid usages for the symmetric key to be generated. Previously, this list was missing the `verify` usage. As a consequence, using the `key-generate --usage verify` option returned an error message. The code has been fixed, and now the `verify` option works as expected. (BZ#1238684)

Subsequent PKI installation no longer fails

Previously, when installing multiple public key infrastructure (PKI) instances in batch mode, the installation script did not wait until the CA instance was restarted. As a consequence, the installation of subsequent PKI instances could fail. The script has been updated and now waits until the new subsystem is ready to handle requests before it continues. (BZ#1446364)

Two-step subordinate CA installation in FIPS mode no longer fails

Previously, a bug in subordinate CA installation in FIPS mode caused two-step installations to fail because the installer required the instance to not exist in the second step. This update changes the workflow so that the first step (installation) requires the instance to not exist, and the second step (configuration) requires the instance to exist.

Two new options, `--skip-configuration`` and `--skip-installation`, have been added to the `pkispawn` command to replace the previous `pki_skip_configuration` and `pki_skip_installation` deployment parameters. This allows you to use the same deployment configuration file for both steps without modifications. (BZ#1454450)

The audit log no longer records success when a certificate request was rejected or canceled

Previously when a certificate request was rejected or canceled, the server generated a `CERT_REQUEST_PROCESSED` audit log entry with `Outcome=Success`. This was incorrect because there was no certificate issued for the request. This bug has been fixed, and the `CERT_REQUEST_PROCESSED` audit log entry for a rejected or canceled request now reads `Outcome=Failure`. (BZ#1452250)

PKI subsystems which failed self tests are now automatically re-enabled on startup

Previously, if a PKI subsystem failed to start due to self test failure, it was automatically disabled to prevent it from running in an inconsistent state. The administrator was expected to re-enable the subsystem manually using `pki-server subsystem-enable` after fixing the problem. However, this was not clearly communicated, potentially causing confusion among administrators who were not always aware of this requirement.

To alleviate this problem, all PKI subsystems are now re-enabled automatically on startup by default. If a self-test fails, the subsystem is disabled as before, but it will no longer require manual re-enabling.

This behavior is controlled by a new boolean option in the `/etc/pki/pki.conf` file, `PKI_SERVER_AUTO_ENABLE_SUBSYSTEMS`. (BZ#1454471)

CERT_REQUEST_PROCESSED audit log entries now include certificate serial number instead of encoded data

Previously, **CERT_REQUEST_PROCESSED** audit log entries included Base64-encoded certificate data. For example:

```
[AuditEvent=CERT_REQUEST_PROCESSED]...[InfoName=certificate]
[InfoValue=MIIDBD...]
```

This information was not very useful because the certificate data would have to be decoded separately. The code has been changed to include the certificate serial number directly into the log entry, as shown in the following example:

```
[AuditEvent=CERT_REQUEST_PROCESSED]...[CertSerialNum=7]
```

(BZ#[1452344](#))

Updating the LDAPProfileSubsystem profile now supports removing attributes

Previously, when updating the **LDAPProfileSubsystem** profile on PKI Server, attributes could not be removed. As a result, PKI Server was unable to load the profile or issue certificates after updating the profile in certain situations. A patch has been applied, and now PKI Server clears the existing profile configuration before loading the new configuration. As a result, updates in the **LDAPProfileSubsystem** profile can now remove configuration attributes. (BZ# [1445088](#))

CHAPTER 24. CLUSTERING

Pacemaker Remote may shut down, even if its connection to the cluster is unmanaged

Previously, if a Pacemaker Remote connection was unmanaged, the Pacemaker Remote daemon would never receive a shutdown acknowledgment from the cluster. As a result, Pacemaker Remote would be unable to shut down. With this fix, if a Pacemaker Remote connection is unmanaged, the cluster now immediately sends a shutdown acknowledgment to Pacemaker Remote nodes that request shutdown, rather than wait for resources to stop. As a result, Pacemaker Remote may shut down, even if its connection to the cluster is unmanaged. (BZ#1388489)

pcs now validates the name and the host of a remote and guest node

Previously, the `pcs` command did not validate whether the name or the host of a remote or guest node conflicted with a resource ID or with a cluster node, a situation that would cause the cluster not to work correctly. With this fix, validation has been added to the relevant commands and `pcs` does not allow a user to configure a cluster with a conflicting name or conflicting host of a remote or guest node. (BZ#1386114)

New syntax of master option in pcs resource create command allows correct creation of meta attributes

Previously, when a `pcs resource create` command included the `--master` flag, all options after the keyword `meta` were interpreted as master `meta` attributes. This made it impossible to create `meta` attributes for the primitive when the `--master` flag was specified. This fix provides a new syntax for specifying a resource as a master slave clone by using the following format for the command:

```
pcs resource create resource_id standard:provider:type|type [resource
options] master [master_options...]
```

This allows you to specify `meta` options as follows:

```
pcs resource create resource_id standard:provider:type|type
[resource_options] meta meta_options... master [master_options...]
```

Additionally, with this fix, you specify a clone resource with the `clone` option rather than the `--clone` flag, as in previous releases. The new format for specifying a clone resource is as follows:

```
pcs resource create resource_id standard:provider:type|type
[resource_options] clone
```

(BZ#1378107)

CHAPTER 25. COMPILER AND TOOLS

The PCRE library now correctly recognizes non-ASCII printable characters as required by Unicode

When matching a Unicode string with non-ASCII printable characters using the Perl Compatible Regular Expressions (PCRE) library, the library was previously unable to correctly recognize printable non-ASCII characters. A patch has been applied, and the PCRE library now recognizes printable non-ASCII characters in UTF-8 mode. (BZ#[1400267](#))

Applications using Bundler to manage dependencies can now properly load the JSON library

Previously, when Bundler was used to manage Ruby application dependencies, it was sometimes impossible to load the JSON library. Consequently, the application failed with a `LoadError`. This caused problems especially because Ruby on Rails no longer explicitly specifies dependency on the JSON library. With this update, JSON is always available on the load path, and the described problem no longer occurs. (BZ#[1308992](#))

git can now be used with HTTP or HTTPS and SSO

Since `libcurl` version 7.21.7, a new parameter for delegating Kerberos tickets is required because of CVE-2011-2192. Previously, Git did not provide a way to set such a parameter. As a consequence, using Git with Single Sign-On on HTTP or HTTPS connections failed. With this update, Git provides a new `http.delegation` configuration variable, which corresponds to the `cURL --delegation` parameter. Users need to set this parameter when delegation of Kerberos tickets is required. (BZ#[1369173](#))

`rescan-scsi-bus.sh --luns=1` now scans only LUNs numbered with 1

The `sg3_utils` package contains utilities that send SCSI commands to devices. In version 1.28-5 and all previous versions of `sg3_utils`, the `rescan-scsi-bus.sh --luns=1` command rescanned only Logical Unit Numbers (LUNs) numbered with 1. After the update to version 1.28-6, `rescan-scsi-bus.sh --luns=1` incorrectly rescanned all LUNs. With this update, the underlying source code has been fixed, and `rescan-scsi-bus.sh --luns=1` now scans only LUNs numbered with 1. (BZ#[1380744](#))

`ps` no longer removes prefixes from wait channel names

The `ps` utility was previously removing the `sys_` and `do_` prefixes from wait channel (`WCHAN`) data. This prevented the user from distinguishing functions with names intentionally containing these prefixes in a `ps` output. The code for prefix removing has been removed, and `ps` now shows full wait channel names. (BZ#[1373246](#))

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

Previously, if the `.history` file was located on a network file system, such as NFS or Samba, the `tcsh` command language interpreter sometimes became unresponsive during the login process. A patch has been applied to avoid `.history` file-locking if `.history` is located on a network file system, and `tcsh` no longer becomes unresponsive in the described situation.

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

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

■

The **lock** option must be the third parameter of the **savehist** option to force **tcsh** to use file-locking when **.history** is located on a network file system. Red Hat does not guarantee that using the **lock** parameter prevents **tcsh** from becoming unresponsive during the login process. (BZ#1388426)

fcoeadm --target no longer causes fcoeadm to crash

Previously, executing the **fcoeadm --target** command sometimes caused the **fcoeadm** utility to terminate unexpectedly with a segmentation fault. With this update, **fcoeadm** has been modified to ignore sysfs paths for non-FCoE targets, and **fcoeadm --target** no longer causes **fcoeadm** to crash. (BZ#1384707)

tar option --directory no longer ignored

Previously, the **--directory** option of the **tar** command was ignored when used in combination with the **--remove-files** option. As a consequence, files in the current working directory were removed instead of the files in the directory specified by the **--directory** option. To fix this bug, new functions and an attribute that retrieve, store, and act upon the **--directory** option have been added. As a result, files are now correctly removed from the directory specified by the **--directory** option. (BZ#1319820)

tar options --xattrs-exclude and --xattrs-include no longer ignored

Previously, the **tar** command ignored the **--xattrs-exclude** and **--xattrs-include** options. To fix this bug, **tar** has been modified to apply include and exclude masks when fetching extended attributes. As a result, the **--xattrs-exclude** and **--xattrs-include** options are no longer ignored. (BZ#1341786)

tar now restores incremental backup correctly

Previously, the **tar** command did not restore incremental backup correctly. Consequently, a file removed in the incremental backup was not removed when restoring. The bug has been fixed, and **tar** now restores incremental backup correctly. (BZ#1184697)

The perl-homedir profile scripts now support csh

Previously, the **perl-homedir** profile scripts were unable to handle the **C shell** (**csh**) syntax. Consequently, when the **perl-homedir** package was installed and the **/etc/sysconfig/perl-homedir** file contained the **PERL_HOMEDIR=0** line, executing the profile scripts resulted in the following error:

```
PERL_HOMEDIR=0: Command not found.
```

This update adds support for the **csh** syntax, and the described problem no longer occurs. (BZ#1122993)

getaddrinfo no longer accessing uninitialised data

On systems with the **nscd** daemon enabled, the **getaddrinfo()** function in the **glibc** library could access uninitialized data and consequently could return false address information. This update prevents uninitialized data access and ensures that correct addresses are returned. (BZ#1324568)

Additional security checks performed in malloc implementation in glibc

Previously, because the **glibc** library was compiled without assertions, the functions implementing **malloc** did not check the heap consistency. This increased the risk that heap-based buffer overflows could be exploited. The heap consistency check has been converted from an assertion into an explicit check. As a result, the security of calls to **malloc** implementation in **glibc** is now increased. (BZ#1326739)

chrpath rebased to version 0.16

The `chrpath` package has been upgraded to upstream version 0.16, which provides a number of bug fixes over the previous version. Notably, the `chrpath` tool could only modify the run path property of 64-bit binaries on 64-bit systems, and 32-bit binaries on 32-bit systems. This bug has been fixed, and `chrpath` on a 64-bit system can now modify run paths of binaries for a 32-bit system and binaries for a 32-bit system on a 64-bit system. (BZ#[1271380](#))

Updated translations for the system-config-language package

To resolve missing translations for `system-config-language`, the following 10 languages have been added; `de`, `es`, `fr`, `it`, `ja`, `ko`, `pt_BR`, `ru`, `zh_CN`, `zh_TW`. (BZ#[1304223](#))

Mutt no longer send emails with an incomplete From header when a host name lacks the domain part

Previously, when a host name did not include the domain name, the `Mutt` email client sent an email with a `From` header that was missing the host name. As a consequence, it was impossible to reply to such an email. This bug has been fixed, and `Mutt` now correctly handles host names that do not contain the domain part. (BZ#[1388512](#))

strace displays correctly the O_TMPFILE flag and mode for open() function

Previously, the `strace` utility did not recognize existence of the `O_TMPFILE` flag for the system function `open()` and its requirement for presence of mode option. As a consequence, the `strace` output did not show name of the respective flag and lacked the mode option value. The `strace` utility has been extended to recognize this situation. As a result, the `O_TMPFILE` flag and mode are displayed correctly. (BZ#[1377847](#))

ld no longer enters an infinite loop when linking large programs

In large programs for the IBM Power Systems architecture the `.text` segment is serviced by two stub sections. Previously, the `ld` linker sizing termination condition was never satisfied when sizing such segments because one of the sections always had to grow. As a consequence, `ld` entered an infinite loop and had to be terminated. `ld` has been extended to recognize this situation and alter the sizing termination condition. As a result, `ld` terminates correctly. (BZ#[1406498](#))

gold warning messages for cross object references to hidden symbols fixed

The `gold` linker produces a warning message when linking shared libraries where the code in one library references a hidden symbol in a second library or object file. Previously, `gold` produced this warning message even if another library or object file provided a visible definition of the same symbol. To fix this bug, `gold` has been extended with a check for this specific case and produces the warning message only if there are no visible definitions of the symbol. As a result, `gold` no longer displays the wrong warning message. (BZ#[1326710](#))

OProfile default event on Intel Xeon® C3xxx Processors with Denverton SOC fixed

Previously, incorrect values were used in the default cycle counting event for `OProfile` on Intel Xeon® C3xxx Processors with Denverton SOC. As a consequence, `OProfile` sampling and counting using the default event did not work. The relevant `OProfile` setting was corrected. As a result, the default event now works on Intel Xeon® C3xxx Processors with Denverton SOC. (BZ#[1380809](#))

CHAPTER 26. DESKTOP

Empathy now can validate certificate chains for Google Talk

Previously, the **Empathy** instant messaging client was not able to validate the certificate chain for Google Talk by ignoring disabled legacy certificate authorities, such as Equifax Secure Certificate Authority, in the chain. As a consequence, **Empathy** prompted the user about an invalid certificate when connecting to Google Talk, even though there was no problem with the chain. This update fixes the bug and **Empathy** now ignores disabled legacy certificate authorities in the list returned by the server and attempts to construct alternate chains that might be valid. (BZ#1386616)

CHAPTER 27. FILE SYSTEMS

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

When starting the `autofs` utility, the `sss` map source was previously sometimes not ready to provide map information, but `sss` did not return an appropriate error to distinguish between the `map does not exist` and `not available` condition. As a consequence, automounting did not work correctly, and `autofs` started without mounts from SSSD. To fix this bug, `autofs` retries asking SSSD for the master map when the `map does not exist` error occurs for a configurable amount of time. Now, you can set the retry timeout to a suitable value so that the master map is read and `autofs` starts as expected. (BZ#1101782)

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

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

- The `README.autofs-schema` file has been added to describe the problem and recommend which schema to use, if possible.
- The schema included in the `autofs` package has been updated to `samples/autofs.schema.new`. (BZ#1383910)

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

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

If the NIS maps are still not available after the configured wait time, the `autofs` configuration `master_wait` option might need to be increased. In the majority of cases, the wait time used by the package is sufficient. (BZ#1383194)

Checking local mount availability with autofs no longer leads to a lengthy timeout before failing

Previously, a server availability probe was not done for mount requests that `autofs` considered local because a bind mount on the local machine is expected to be available for use. If the bind mount failed, an NFS mount on the local machine was then tried. However, if the NFS server was not running on the local machine, the mount attempt sometimes suffered a lengthy timeout before failing.

An availability probe has been added to the case where a bind mount is first tried, but fails, and `autofs` now falls back to trying to use an NFS server on the local machine. As a result, if a bind mount on the local machine fails, the fallback to trying an NFS mount on the local machine fails quickly if the local NFS server is not running. (BZ#1420574)

The journal is marked as idle when mounting a GFS2 file system as read-only

Previously, the kernel did not mark the file system journal as idle when mounting a GFS2 file system as read-only. As a consequence, the `gfs2_log_flush()` function incorrectly tried to write a header block to the journal and a sequence-out-of-order error was logged. A patch has been applied to mark the journal idle when mounting a GFS2 file system as read-only. As a result, the mentioned error no longer occurs in the described scenario. (BZ#1213119)

The `id` command no longer shows incorrect UIDs and GIDs

When running Red Hat Enterprise Linux on an NFSv4 client connected to an NFSv4 server, the `id` command showed incorrect UIDs and GIDs after the key expired out of the NFS idmapper keyring. The problem persisted for 5 minutes, until the expired keys were garbage collected, after which the new key was created in the keyring and the `id` command provided the correct output. With this update, the keyring facility has been fixed, and the `id` command no longer shows incorrect output under the described circumstances. (BZ#1408330)

Labeled NFS is now turned off by default

The SELinux labels on a Red hat Enterprise Linux NFS server are not normally visible to NFS clients. Instead, NFS clients see all files labeled as type `nfs_t` regardless of what label the files have on the server.

Since Red Hat Enterprise Linux 7.3, the NFS server has the ability to communicate individual file labels to clients. Sufficiently recent clients, such as recent Fedora clients, see NFS files labeled with the same labels that those files have on the server. This is useful in certain cases, but it can also lead to unexpected access permission problems on recent clients after a server is upgraded to Red Hat Enterprise Linux 7.3 and later.

Note that labeled NFS support is turned off by default on the NFS server. You can re-enable labeled NFS support by using the `security_label` export option. (BZ#1406885)

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

If an `autofs` mount reached a shutdown state, and a mount request arrived and was processed before the mount-handling thread read the shutdown notification, the mount-handling thread previously exited without cleaning up the `autofs` mount. As a consequence, the main program never reached its exit condition and entered an infinite loop, as the `autofs`-managed mount was left mounted. To fix this bug, the exit condition check now takes place after each request is processed, and cleanup operations are now performed if an `autofs` mount has reached its shutdown state. As a result, the `autofs` daemon now exits as expected at shutdown. (BZ#1420584)

`autofs` is now more reliable when handling namespaces

Previously, the `autofs` kernel module was unable to check whether the last component of a path was a mount point in the current namespace, only whether it was a mount point in any namespace. Due to this bug, `autofs` sometimes incorrectly decided whether a mount point cloned into a propagation private namespace was already present.

As a consequence, the automount point failed to be mounted and the error message `Too many levels of symbolic links` was returned. This happened, for example, when a `systemd` service that used the `PrivateTmp` option was restarted while an `autofs` mount was active.

With this update, a namespace-aware mounted check has been added in the kernel. As a result, `autofs` is now more resilient to cases where a mount namespace that includes `autofs` mounts has been cloned to a propagation private namespace.

For more details, see the KBase article at <https://access.redhat.com/articles/3104671>. (BZ#1320588)

CHAPTER 28. INSTALLATION AND BOOTING

Automatic partitioning now works when installing on a single FBA DASD on IBM z Series

Previously, when installing Red Hat Enterprise Linux 7 on IBM z Series systems with a single Fixed Block Architecture (FBA) Direct Access Storage Device (DASD) with the `cms` disk layout as the target, automatic partitioning failed because the installer attempted to create multiple partitions on the device, which is not supported on `cms`-formatted FBA DASDs. This caused the installation finish with a corrupted disk.

With this update, the installer first creates a `msdos` partition table on the target DASD, which allows up to three partitions on the device. As long as the installer only creates three or fewer partitions, the installation will succeed. Note that it is recommended to use the `autopart --nohome` Kickstart option to ensure that the installer does not create a separate `/home` partition. (BZ#1214407)

Activation of bridge configured in Kickstart no longer fails when Kickstart proceeds from the disk

Previously, if the bridge device was configured in a Kickstart file and the Kickstart file was fetched from the disk, the lack of network connection meant that the bridge was not created and the installation failed at an early stage. With this update, the bridge Kickstart configuration is passed to the `dracut` tool at an early stage. As a result, `dracut` can create and activate the bridge device even when no network is required at the early stage of installation. (BZ#1373360)

Anaconda now correctly allows creating users without passwords

Previously, it was not possible to deselect the `Require a password to use this account` option in the `Create User` screen during an interactive installation. As a consequence, all user accounts created during the installation required a password. This bug has been fixed, and creating users with no password is now possible. (BZ#1380277)

Minimal installation no longer installs open-vm-tools-desktop and dependencies

The `open-vm-tools-desktop` package was previously marked as default in the `@platform-vmware` package group (Virtualization utilities and drivers for VMWare). This group is automatically installed by `Anaconda` when it detects that the installation is using a `VMWare` hypervisor. At the same time, this package has many dependencies including a large number of X libraries which are not useful in a minimal installation, and this was causing `Anaconda` to install a high number of unnecessary packages.

The `open-vm-tools-desktop` package is now optional in the `@platform-vmware` group, and therefore not being installed by default. The other package in the group, `open-vm-tools`, remains mandatory and is therefore installed by default. (BZ#1408694)

Anaconda no longer generates invalid Kickstart files

Previously, if a Kickstart file was used during an installation which defined some LVM logical volumes absolutely (the `--size=` parameter) and others relatively (the `--percent=` parameter), the resulting Kickstart file which is saved on the installed system, `anaconda-ks.cfg`, defined all logical volumes using both of these parameters. These parameters are mutually exclusive, and the generated Kickstart file was therefore invalid. With this update, `Anaconda` correctly handles usage of relative and absolute sizes, and the resulting post-installation Kickstart files are valid. (BZ#1317370)

Anaconda no longer fails to identify RAID arrays specified by name

Previously, when a RAID array was specified by name in the `ignoredisk` or `clearpart` command in a Kickstart file, the installation could not proceed because RAID names are not available during initial stages of the installation. This update improves RAID support by ensuring that `Anaconda` also checks devices in `/dev/md/` for a matching name. For example, if the Kickstart file contains the command

ignoredisk --only-use=myraid, Anaconda will now also attempt to find an array located at `/dev/md/myraid`. This allows the installer to locate RAID arrays specified by name at any point during the installation, and enables specifying only RAID array names in Kickstart files. (BZ#1327439)

Kickstart no longer accepts passwords that are too short

Previously, when using a Kickstart file to install Red Hat Enterprise Linux 7, the Anaconda installer immediately accepted passwords shorter than the minimal length defined by the `--minlen` Kickstart option, if the password was sufficiently strong (quality value 50 or above by default). This bug has been fixed, and the `--minlen` option now works even with strong passwords. (BZ#1356975)

Initial Setup now correctly opens in a graphical interface over SSH on IBM z Systems

Previously, when connecting to an IBM z Systems machine using SSH, the text version of the **Initial Setup** interface opened even if X forwarding was enabled. This bug has been fixed, and the graphical version of **Initial Setup** now opens correctly when using X forwarding. (BZ#1378082)

Extra time is no longer needed for installation when geolocation services are enabled

When installing Red Hat Enterprise Linux 7.3 with limited or no internet access, the installer previously paused for several minutes in the Installation Summary screen with the Security Policy section being **Not ready**. This was caused by the geolocation service being unable to determine the system's location. Consequently, the installation could not proceed before the service timed out. With this update, the geolocation service correctly times out if it can not find the location within 3 seconds, and the installation can proceed almost immediately even with limited or no network connection. (BZ#1380224)

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

When moving one or more IP aliases from one server to another, associated IP addresses may be unreachable for some time, depending on the Address Resolution Protocol (ARP) time-out value that is configured in the upstream router. This bug has been addressed in the `initscripts` package, and `ifup-aliases` now updates other systems on the network significantly faster in this situation. (BZ#1367554)

The netconsole utility now launches correctly

Previously, if `nameserver` address lines were not present in the `/etc/resolv.conf` file, launching `netconsole` sometimes resulted in an error and `netconsole` did not start. The `initscripts` package has been updated, and `netconsole` now starts correctly in this situation. (BZ#1278521)

rc.debug kernel allows easier debugging of initscripts

This enhancement introduces the `rc.debug` option for the kernel command line. Adding the `rc.debug` option to the kernel command line prior to booting produces a log of all the activity of the `initscripts` files during the boot and termination processes. The log appears as part of the `/var/log/dmesg` log file. As a result, adding the `rc.debug` option to the kernel command line enables easier debugging of `initscripts` if needed. (BZ#1394191)

The system no longer fails to terminate with /usr on iSCSI or NFS

In previous versions of Red Hat Enterprise Linux 7, the termination of the system sometimes failed and the system remained hung if the `/usr` folder was mounted over a network (for example, `NFS` or `iSCSI`). This issue has been resolved, and the system should now shut down normally. (BZ#1369790, BZ#1446171)

rhel-autorelabel no longer corrupts the filesystem

In previous versions of Red Hat Enterprise Linux 7, forcing the SELinux autorelabel by creating the `/.autorelabel` file sometimes partially corrupted the filesystem. This made the system unbootable. A patch has been applied to prevent this behaviour. As a result, applying the `autorelabel` operation using the `touch /.autorelabel` command is no more expected to corrupt the filesystem. (BZ#1385272)

The `rpmbuild` command now correctly processes Perl requires

Previously, a bug in `rpm` caused `my variable = << blocks` to be treated as code instead of string constants when building packages using the `rpmbuild` command. This caused `rpm` to add unintended dependencies to packages being built in cases where the `variable` contained the word `use` followed by another word. With this update, `rpm` correctly skips these blocks when searching for dependencies, and packages no longer contain unintended dependencies. (BZ#1378307)

Installer now correctly recognizes BIOS RAID devices when using `ignoredisk` in Kickstart

Previously, some BIOS RAID devices were not correctly recognized during installation when using a Kickstart file with the `ignoredisk --onlyuse=<bios raid name>` command. This caused the installation to fail and report lack of free space because the device could not be used. With this update, Anaconda recognizes BIOS RAID devices reliably when they are specified in a Kickstart file, and installations no longer fail in these circumstances. (BZ#1327463)

Single quotes now work for values in the `ifcfg-*` files

Previously, it was only possible to specify values by using double quotes in the `ifcfg-*` files. Using single quotes did not work. With this update, single quotes work too, for example:

```
ONBOOT='yes'
```

(BZ#1428574)

`rhel-import-state` no longer changes access permissions for `/dev/shm/`, allowing the system to boot correctly

Previously, problems during the boot-up process occurred due to the introduction of a new script in a `dracut` update. The new script changed the access permission to the `/dev/shm/` directory when the `dracut` utility placed the directory to the `/run/initramfs/state/`. With this update, `rhel-import-state` no longer changes the access permissions for `/dev/shm/`, and the system starts correctly. (BZ#1406254)

Backward compatibility enabled for Red Hat Enterprise Linux 6 `initscripts`

The `initscripts` files in Red Hat Enterprise Linux 7 have been patched to enable backward compatibility and to prevent possible regressions when doing an upgrade from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7. (BZ#1392766)

`initscripts` now specifies `/etc/rwtab` and `/etc/statetab` as configuration files

Previously, a reinstallation of the `initscripts` package replaced the `/etc/rwtab` and `/etc/statetab` files. If these files had contained user's configuration, the reinstallation process overwrote it.

The `initscripts` package has been updated to specify the `/etc/rwtab` and `/etc/statetab` files as configuration files. If these files are modified by the user, performing the reinstallation now creates the `*.rpmnew` files containing the new configuration in the `/etc/` folder. As a result of this update, a reinstallation of the `initscripts` package leaves the `/etc/rwtab` and `/etc/statetab` files intact. (BZ#1434075)

The `ifup` script no longer slows down `NetworkManager`

Previously, the `ifup` script was very slow when notifying `NetworkManager`. This particularly affected Red Hat Virtualization (RHV) network startup times. A patch has been applied to initscripts, and the described problem no longer occurs. (BZ#[1408219](#))

Gnome Initial Setup can now be disabled by the `firstboot --disable` command in kickstart

With this update, the `gnome-initial-setup` package has been fixed to respect the `firstboot --disable` kickstart command. As a result, Gnome Initial Setup can be robustly turned off during a kickstart installation and users are no longer forced to create user account on the first boot under the described circumstances as long as the installation kickstart contains the `firstboot --disable` command. (BZ#[1226819](#))

Setting `NM_CONTROLLED` now works correctly across all the `ifcfg-*` files

When the `NM_CONTROLLED=no` parameter was set for an interface in its `ifcfg-*` file, other interfaces in some cases inherited this configuration. This behaviour prevented the `NetworkManager` daemon from controlling these interfaces. The issue has now been resolved, and setting the `NM_CONTROLLED` parameter now works correctly across all the `ifcfg-*` files. As a result, the user can choose which interface is controlled by `NetworkManager`, and which is not. (BZ#[1374837](#))

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

The `dhclient` command incorrectly sent `localhost` to the DHCP server as the host name when the `hostname` variable was not set. This has been fixed, and `dhclient` no longer sends an incorrect host name in these situations. (BZ#[1398686](#))

The `initscripts` utility now handles LVM2 correctly

Previously, later versions of the `initscripts` utility made use of a new `--ignoresthecluster` option for the `vgchange` command during boot. This option was missing in earlier versions of the `lvm2` utilities. As a consequence, systems using earlier versions of the Logical Volume Manager device mapper (LVM2) could fail to boot correctly. With this update, the `initscripts` RPM indicates the version of `lvm2` required, and if a sufficient version is installed, systems with LVM2 boot correctly. (BZ#[1398683](#))

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

Previously, when a tunnel interface was present, the `service network stop` command incorrectly attempted to stop services which had been stopped already, displaying an error message. This bug has been fixed, and the `service network stop` command now stops only running services. (BZ#[1398679](#))

`ifdown` on a loopback device now works correctly

In previous versions of Red Hat Enterprise Linux 7, executing the `ifdown` command on a local loopback device failed to remove the device. A patch has been applied, and the removal of an existing loopback device using `ifdown` now succeeds. (BZ#[1398678](#))

Scripts in `initscripts` handle static IPv6 address assignment more robustly

Previously, scripts in the `initscripts` package sometimes failed to correctly assign static IPv6 addresses if a Router Advertisement (RA) was received during system initialization. This bug has been fixed, and now the statically assigned address is correctly applied in the described situation. (BZ#[1398671](#))

Deselecting an add-on option in Software Selection no longer requires a double-click

When installing Red Hat Enterprise Linux 7.3, the user had to double-click in order to deselect an add-

on checkbox after a **Base environment** change. The bug occurred in the **Software Selection** dialogue of the graphical installation. With this update, the system no longer requires double-clicking when deselecting an option after a **Base environment** change. A single click is sufficient. (BZ#1404158)

The target system hostname can be configured via installer boot options in Kickstart installations

In Red Hat Enterprise Linux 7.3, the hostname specified via the **Anaconda** installer boot options during a Kickstart installation was previously incorrectly not set for installed system and the default **localhost.localdomain** hostname value was used instead. With this update, **Anaconda** has been fixed to apply the hostname set by the boot option to the target system configuration. As a result, users can now configure the target system hostname via the installer boot options also for Kickstart installations. (BZ#1441337)

Anaconda no longer asks for Installation Source verification after network configuration

Previously, during an **Anaconda** installation from a repository, when the user changed network settings after repository packages had already been selected, the Installation Source required verification. This request was made even when the repository was still reachable after the network change, resulting in an unnecessary step. With this update, the **Anaconda** installer keeps the original source repository and verifies whether it is still reachable after the Network & Hostname configuration. As a result, the user is only required to reconfigure the Installation Source if the original repository is not reachable. (BZ#1358778)

Disks using the OEMDRV label are now correctly ignored during automatic installation

The OEMDRV disk label is used on driver update disks during installation. Due to a bug, disks with this label were being used by **Anaconda** as installation targets during automatic installations, which meant they were being erased and used as part of the installed system storage. This update ensures that **Anaconda** ignores disks with this label unless they are explicitly selected as installation targets, and the problem no longer occurs. (BZ#1412022)

CHAPTER 29. KERNEL

RAID 4 and RAID 10 creation and activation fully supported

In Red Hat Enterprise Linux 7.3, existing RAID4 or RAID10 logical volumes created with previous releases failed to activate. Additionally, users were instructed not to create new RAID4 logical volumes created under Red Hat Enterprise Linux 7.3 as these might fail to activate with later releases and updates. With this update, Red Hat Enterprise Linux 7.4 fully supports RAID 4 and RAID 10 creation and activation and rejects invalid RAID 4 and RAID 10 layouts which may have been created with Red Hat Enterprise Linux 7.3. (BZ#1385149)

kdump now works with legacy type 12 NVDIMMs

Previously, systems with legacy type 12 Non-Volatile Dual In-line Memory Modules (NVDIMMs), either real dual in-line memory modules (DIMMs), or emulated using the `memmap=XG!YG` kernel command line parameter, were unable to successfully capture a kernel crash dump. For systems with real NVDIMMs, attempts to capture a kernel crash dump occasionally resulted in data corruption. With this update, the underlying source code has been fixed, and the systems with legacy type 12 NVDIMMs can now capture the kernel crash dump as expected. (BZ#1351098)

Creating a file that inherits ACLs no longer loses mask

Previously, creating a file that inherited Access Control Lists (ACLs) caused the mask to be lost, unlike on a local file system. With this update, clients using NFSv4.2 can set the `umask` attribute when creating files to cause the server to apply `umask` always except when inheriting permissions from the parent directory. As a result, new NFS files get the same permissions as files created locally. Note that you need to apply this update on both your NFS clients and NFS servers, and mount with the `-overs=4.2` parameter. (BZ#1217546)

CHAPTER 30. REAL-TIME KERNEL

Removing USB no longer causes a `might_sleep()` warning on MRG Realtime kernel

Removing an USB device on MRG Realtime kernel previously caused taking a sleeping spin lock with interruptions disabled. Consequently, the system logged a `might_sleep()` warning. This update replaces the `local_irq_disable` and `local_irq_enable` calls by `local_irq_disable_nort` and `local_irq_enable_nort`, respectively, which fixes this bug. (BZ#1443711)

CHAPTER 31. NETWORKING

SNMP response is no longer timed out

Previously, all the Simple Network Management Protocol version 1 (SNMPv1) and SNMPv2c responses that followed an SNMPv3 message were checked against the last recorded SNMPv3 `max_message_size` property. As a consequence, an SNMPv3 request with a small `max_message_size` could lead to SNMPv1 and SNMPv2c bulk requests timing out. With this update, the session maximum message size is checked only for SNMPv3 requests, and the SNMPv1 and SNMPv2c response is no longer timed out. (BZ#1324306)

ICMP redirects no longer cause kernel to crash

Previously, a socket failed to be locked between user space and the process of Internet Control Message Protocol (ICMP) redirect packets, creating a race condition. As a consequence, kernel terminated unexpectedly. The bug has been fixed by skipping the process of ICMP redirect packets when the socket is locked by user space and now the described problem no longer occurs. (BZ#1387485)

The `net.ipv4.ip_nonlocal_bind` kernel parameter is set in name spaces

Previously, using a floating IP address inside a network name space in some cases failed with the following error message:

```
bind: Cannot assign requested address.
```

With this update, the kernel respects setting of the `net.ipv4.ip_nonlocal_bind` parameter to 1 in name spaces, and the floating IP address is now assigned as expected. (BZ#1363661)

The netfilter REJECT rule now works on SCTP packets

Previously, the `conntrack` tool did not check the `CRC32c` value for Stream Control Transmission Protocol (SCTP) packets. As a consequence, the netfilter `REJECT` rule was not applied as expected on SCTP packets. The bug has been fixed by setting `CHECKSUM_UNNECESSARY` on SCTP packets which have valid `CRC32c`. As a result, the netfilter `REJECT` is allowed to generate an Internet Control Message Protocol (ICMP) response. (BZ#1353218)

NetworkManager no longer duplicates a connection with already-set `DHCP_HOSTNAME`

Previously, after a restart of the `NetworkManager` service, a connection with an already-set `DHCP_HOSTNAME` property was duplicated. Consequently, a DHCP lease was not always renewed upon its expiry. With this update, the connection is no longer duplicated, and a DHCP lease is correctly renewed in this scenario.

Note that the fix includes ignoring the already-set hostname properties in the matching process. To avoid possible problems, remove all unused connections with an incorrect `ipv4.dhcp-hostname`. For more information, see <https://access.redhat.com/articles/2948041>. (BZ#1393997)

Improved SCTP `congestion_window` management

Previously, small data chunks caused the Stream Control Transmission Protocol (SCTP) to account the `receiver_window` (`rwnd`) values incorrectly when recovering from a `zero-window situation`. As a consequence, window updates were not sent to the peer, and an artificial growth of `rwnd` could lead to packet drops. This update properly accounts such small data chunks and ignores the `rwnd` pressure values when reopening a window. As a result, window updates are now sent, and the announced `rwnd` reflects better the real state of the receive buffer. (BZ#1084802)

Value of DCTCP `alpha` now drops to 0 and `cwnd` remains at values more than 137

Previously, the `alpha` value of Datacenter TCP (DCTCP) was shifted before subtraction, causing

precision loss. As a consequence, the real `alpha` value did not fall below 15 and uncongested flows eventually dropped to a `congestion_window (cwnd)` value of 137. This bug has been fixed by canceling the shift operation when `alpha` is low. As a result, `alpha` drops to 0 and `cwnd` remains at values more than 137 for uncongested flows. (BZ#1370638)

ss now displays correctly cwnd

Previously, the `ss` utility displayed Transmission Control Protocol congestion window (TCP `cwnd`) values from the kernel, performing a cast from unsigned to signed 32-bit integer. As a consequence, some values can overflow and be interpreted as a negative value. With this update, the `ss` code has been fixed, and the utility no longer displays negative `cwnd` values. (BZ#1375215)

Value of cwnd no longer increases using DCTCP

Previously, the `congestion_window (cwnd)` increased unexpectedly after a packet loss. As a consequence, the Data Center TCP (DCTCP) congestion control module became ineffective in avoiding congestion, because repeated problems on the same flow occurred. With this update, the `cwnd` value is saved on loss and the old one is restored on recovery. As a result, `cwnd` remains stable. (BZ#1386923)

Negated range matches have been fixed

Previously, using a range of values in a negated match would never evaluate as true. With this update, such matches work as expected. For example:

```
# nft add rule ip ip_table filter_chain_input ip length != 100-200 drop
```

now correctly drops packets smaller than 100 bytes or larger than 200 bytes. (BZ#1418967)

The nmcli connection show command now displays the correct output for both empty and NULL values

Previously, the output of the `nmcli connection show` command did not display consistently the `empty` and `NULL` values among different properties. As a consequence, the `empty` values were displayed by `--` or without a value. With this update, the output of the `nmcli connection show` command displays `--` for both `empty` and `NULL` values in `normal` or `pretty` modes.

Note that in `terse` mode, values are printed only in their raw form and the `empty` and `NULL` values are not printed at all. (BZ#1391170)

snmpd no longer rejects large packets from AgentX subagents

Previously, the SNMP daemon (`snmpd`) limited the size of packets sent from AgentX subagents to 1472 bytes. This caused `snmpd` to refuse large packets from AgentX subagents. The packet size limit has been increased to 65535 bytes. As a result, `snmpd` no longer rejects large packets from AgentX subagents. (BZ#1286693)

Macvlan can now be unregistered correctly

Previously, attempts to unregister the `Macvlan` driver failed with broken `sysfs` links from or to devices in another namespace. With this update, `Macvlan` has been fixed, thus fixing this bug. (BZ#1412898)

CHAPTER 32. SECURITY

Configurations that depend on chrooting in user-non-searchable paths now work properly

In Red Hat Enterprise Linux 7.3, the `chroot` process in the `OpenSSH` tool had been changed to help harden the SELinux system policy, and root UID was dropped before performing `chroot`. Consequently, existing configurations that depend on chrooting in user-non-searchable paths stopped working. With this update of the `openssh` packages, the change has been reverted. Additionally, the problem has been fixed in the SELinux system policy by allowing confined users to use `OpenSSH chroot` if the administrator enables the `selinuxuser_use_ssh_chroot` boolean. The described configurations now work in the same way as in Red Hat Enterprise Linux 7.2. (BZ#[1418062](#))

firewalld now supports all ICMP types

Previously, the Internet Control Message Protocol (ICMP) type list was not complete. As a consequence, some ICMP types such as `packet-too-big` could not be blocked or allowed. With this update, support for additional ICMP types has been added, and the `firewalld` service daemon now allows to handle all ICMP types. (BZ#[1401978](#))

docker.pp replaced with container.pp in selinux-policy

Prior to this update, the `container.te` file in the `container-selinux` package contained Docker interfaces, which point to the equivalent container interfaces, and also the `docker.if` file. Consequently, when compiling the `container.te` file, the compiler warned about duplicate interfaces. With this update, the `docker.pp` file in the `selinux-policy` package has been replaced with the `container.pp` file, and the warning no longer occurs in the described scenario. (BZ# [1386916](#))

Recently-added kernel classes and permission defined in selinux-policy

Previously, several new classes and permissions had been added to the kernel. As a consequence, these classes and permissions that were not defined in the system policy caused SELinux denials or warnings. With this update, all recently-added kernel classes and permissions have been defined in the `selinux-policy` package, and the denials and warnings no longer occur. (BZ#[1368057](#))

nss now properly handles PKCS#12 files

Previously, when using the `pk12util` tool to list certificates in a PKCS#12 file with strong ciphers using PKCS#5 v2.0 format, there was no output. Additionally, when using `pk12util` to list certificates in a PKCS#12 file with the SHA-2 Message Authentication Code (MAC), a MAC error was reported, but no certificates were printed. With this update, importing and exporting PKCS#12 files has been changed to be compatible with the `OpenSSL` handling, and PKCS#12 files are now processed properly in the described scenarios. (BZ#[1220573](#))

OpenSCAP now produces only useful messages and warnings

Previously, default scan output settings have been changed, and debug messages were also printed to standard output. As a consequence, the `OpenSCAP` output was full of errors and warnings. The output was hard to read and the `SCAP Workbench` was unable to handle those messages, too. With this update, the change of default output setting has been reverted, and `OpenSCAP` now produces useful output. (BZ#[1447341](#))

AIDE now logs in the syslog format

With this update, the `AIDE` detection system with the `syslog_format` option logs in the `rsyslog`-compatible format. Multiline logs cause problems while parsing on the remote `rsyslog` server. With the new `syslog_format` option, `AIDE` is now able to log with every change logged as a single line. (BZ#[1377215](#))

Installations with the OpenSCAP security-hardening profile now proceed

Prior to this update, typos in the scap-security-guide package caused the **Anaconda** installation program to exit and restart a machine. Consequently, it was not possible to select any of the security-hardened profiles such as Criminal Justice Information Services (CJIS) during the Red Hat Enterprise Linux 7.4 installation process. The typos have been fixed, and installations with the **OpenSCAP** security-hardening profile now proceed. (BZ#1450731)

OpenSCAP and SSG are now able to scan RHV-H systems correctly

Previously, using the OpenSCAP and SCAP Security Guide (SSG) tools to scan a Red Hat Enterprise Linux system working as a Red Hat Virtualization Host (RHV-H) returned **Not Applicable** results. With this update, OpenSCAP and SSG correctly identify RHV-H as Red Hat Enterprise Linux, which enables OpenSCAP and SSG to scan RHV-H systems properly. (BZ#1420038)

OpenSCAP now handles also uncompressed XML files in a CVE OVAL feed

Previously, the **OpenSCAP** tool was able to handle only compressed CVE OVAL files from a feed. As a consequence, the CVE OVAL feed provided by Red Hat cannot be used as a base for vulnerability scanning. With this update, **OpenSCAP** supports not only ZIP and BZIP2 files but also uncompressed XML files in a CVE OVAL feed, and the CVE OVAL-based scanning works properly without additional steps. (BZ#1440192)

CHAPTER 33. SERVERS AND SERVICES

rear now correctly preserves Linux capabilities

Previously, `rear` did not preserve the Linux capabilities set on the original system during the backup phase. The recovered system was subsequently missing these capabilities. With this update, Linux capabilities are correctly preserved if the directive `NETFS_RESTORE_CAPABILITIES` is set to the `y` option in the `/usr/share/rear/conf/default.conf` configuration file. (BZ#1343119)

sblim-cmpi-fsvol no longer shows file systems mounted with DM as disabled

Previously, the `sblim-cmpi-fsvol` Common Information Model (CIM) provider could not identify the file systems (FS) that were mounted with the Device Mapper (DM) correctly. Consequently, when enumerating the `CIM_UnixLocalFileSystem` class instances, `sblim-cmpi-fsvol` showed some FS that were already mounted as disabled. With this update, `sblim-cmpi-fsvol` has been fixed to parse the output of the `dmsetup` command instead of parsing the output of the `mount` command for FS mounted with DM. As a result, `sblim-cmpi-fsvol` now shows the FS mounted with DM correctly. (BZ#1136116)

SPNEGO in Cyrus SASL is now compatible with Microsoft Windows

Prior to this update, the Red Hat Enterprise Linux implementation of Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO) in Cyrus Simple Authentication and Security Layer (SASL) was not compatible with the Microsoft Windows counterpart. As a consequence, the Red Hat Enterprise Linux tools using the `cyrus-sasl` packages were not able to use SPNEGO when attempting to connect to Windows services. These tools were also not able to accept connections from Windows clients. The `cyrus-sasl` packages have been fixed, and SPNEGO in the Red Hat Enterprise Linux Cyrus SASL version is now compatible with the Microsoft Windows counterpart. (BZ#1421663)

Data are no longer lost when the MariaDB init script fails

Previously, if the `MariaDB` init script failed, it called `rm -rf` on the whole directory. This could consequently lead to a loss of data or even deletion of the mount point. With this update, several additional checking mechanisms have been added to the init script. Now, if the script fails, it removes only files newer than a timestamp generated prior to the critical file operations. In addition, a set of human-readable status reports and error messages have been added. (BZ#1356897)

ypbind no longer starts before access to the network is guaranteed

The `ypbind` service was set to start after the `systemd` target `network.target`. However, `network.target` does not guarantee the networking functionality required by `ypbind`. As a consequence, the `ypbind` service sometimes did not have access to the network when started during the boot process. The service file for `ypbind` has been changed to start `ypbind` after the target `network-online.target`, and `ypbind` is now guaranteed to have access to the network when started. (BZ#1382804)

Remote users' account settings are no longer reverted to default settings on restart due to ypbind

Because of wrong service starting order, `ypbind` did not start before all Name Service Switch (NSS) look-up operations had been completed. This caused the user's account settings file to revert to default settings on restart for users fulfilling all of the following conditions:

- Using Gnome Display Manager auto-login
- Using NIS authentication
- Home directory located on NFS

The `ypbind` service file ordering has been fixed to make `ypbind` start before the user/group database is set up. The user's account settings file is now handled properly. (BZ#1217435)

ypasswd no longer crashes due to Network Information System security features used

The `ypasswd` client tried to use a wrong string as salt when checking passwords because it did not recognize the following situations:

- The NIS server was configured to use the `passwd.adjunct` map
- The variable `MERGE_PASSWD=false` was set in the NIS server's file `/var/yp/Makefile`

As a consequence, `ypasswd` failed with the following error message: `crypt() failed`. The `ypasswd` client has been fixed to recognize these situations and now delegates the checks to the `ypasswdd` daemon running on the server. (BZ# [1401432](#))

Evince now displays PostScript files again

Due to a bug, the `evince` document viewer failed to display the content of PostScript files. A patch has been applied, and `evince` now displays PostScript files again. (BZ# [1411725](#))

db_verify no longer causes libdb to run out of free mutexes

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

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

leaving the database in an inconsistent state. This bug has been fixed, `libdb` now correctly releases mutexes, and the described problem no longer occurs. (BZ#[1277887](#))

ghostscript no longer becomes unresponsive in some situations

Under certain circumstances, the `ghostscript` application previously entered an infinite loop, became unresponsive, and caused excessive CPU load. This update fixes the underlying code, which prevents the described problem from occurring. (BZ#[1424752](#))

Converting postscript to PDF no longer causes ps2pdf to terminate unexpectedly

Previously, converting a postscript file to PDF in some cases caused the `ps2pdf` utility to terminate unexpectedly with a segmentation fault. This bug has been fixed, and converting postscript to PDF no longer causes `ps2pdf` to crash. (BZ#[1390847](#))

sapconf now works correctly with higher kernel.shmall and kernel.shmmax values

Previously, the `kernel.shmall` and `kernel.shmmax` values were increased by default displaying an error in the `sapconf` utility. Consequently, `sapconf` failed with the following error message:

```
integer expression expected
```

This update adds a new check which allows the high values of `kernel.shmall` and `kernel.shmmax`, and the described problem no longer occurs. (BZ#[1391881](#))

CHAPTER 34. STORAGE

lvconvert --repair now works properly on cache logical volumes

Due to a regression in the `lvm2-2.02.166-1.el` package, released in Red Hat Enterprise Linux 7.3, the `lvconvert --repair` command could not be run properly on cache logical volumes. As a consequence, the `Cannot convert internal LV` error occurred. The underlying source code has been modified to fix this bug, and `lvconvert --repair` now works as expected. (BZ# [1380532](#))

LVM2 library incompatibilities no longer cause device monitoring to fail and be lost during an upgrade

Due to a bug in the `lvm2-2.02.166-1.el` package, released in Red Hat Enterprise Linux 7.3, the library was incompatible with earlier versions of Red Hat Enterprise Linux 7. The incompatibility could cause device monitoring to fail and be lost during an upgrade. As a consequence, device failures could go unnoticed (RAID), or out-of-space conditions were not handled properly (thin-p). This update fixes the incompatibility, and the logical volume monitoring now works as expected. (BZ#[1382688](#))

be2iscsi driver errors no longer cause the system to become unresponsive

Previously, the operating system sometimes became unresponsive due to `be2iscsi` driver errors. This update fixes `be2iscsi`, and the operating system no longer hangs due to `be2iscsi` errors. (BZ#[1324918](#))

Interaction problems no longer occur with the lvm2 daemon when mirror segment type is used

Previously, when the legacy `mirror` segment type was used to create mirrored logical volumes with 3 or more legs, interaction problems could occur with the `lvm2` daemon. Problems observed occurred only after a second device failure, when mirror fault policies were set to the non-default `allocate` option, when `lvm2` was used, and there had been no reboot of the machine between device failure events. This bug has been fixed, and the described interaction problems no longer occur. (BZ#[1380521](#))

The multipathd daemon no longer shows incorrect error messages for blacklisted devices

Previously, the `multipathd` daemon showed incorrect error messages that it could not find devices that were blacklisted, causing users to see error messages when nothing was wrong. With this fix, `multipath` checks if a device has been blacklisted before issuing an error message. (BZ#[1403552](#))

Multipath now flags device reloads when there are no usable paths

Previously, when the last path device to a multipath device was removed, the state of `lvm2` became incorrect, and `lvm` devices on top of multipath could stop working correctly. This was because there was no way for the device-mapper to know how many working paths there were when a multipath device got reloaded. Because of this, the multipath `udev` rules that dealt with disabling scanning and other `dm` rules only worked when multipath devices lost their last usable path because it failed, instead of because of a device table reload. With this fix, multipath now flags device reloads when there are no usable paths, and the multipath `udev` rules now correctly disable scanning and other `dm` rules whenever a multipath device loses its last usable path. As a result, the state of `lvm2` remains correct, and `LVM` devices on top of multipath continue working correctly. (BZ#[1239173](#))

Read requests sent after failed writes will always return the same data on multipath devices

Previously, if a write request was hung in the `rbd` module, and the iSCSI initiator and multipath layer decided to fail the request to the application, read requests sent after the failure may not have reflected the state of the write. This was because When a Ceph `rbd` image is exported through multiple iSCSI targets, the `rbd` kernel module will grab the exclusive lock when it receives a write request. With

this fix, The `rbld` module will grab the exclusive lock for both reads and writes. This will cause hung writes to be flushed and or failed before executing reads. As a result, read requests sent after failed writes will always return the same data. (BZ#[1380602](#))

When a path device in a multipath device switches to read-only, the multipath device will be reloaded read-only

Previously, when reloading a multipath device, the multipath code always tried to reload the device read-write first, and then failed back to read-only. If a path device had already been opened read-write in the kernel, it would continue to be opened read-write even if the device had switched to read-only mode and the read-write reload would succeed. Consequently, when path devices switched to from read-write to read-only, the multipath device would still be read-write (even though all writes to the read-only device would fail). With this fix, `multipathd` now reloads the multipath device read-only when it gets a `uevent` showing that a path devices has become read-only. As a result, when a path device in a multipath device switches to read-only, the multipath device will be reloaded read-only. (BZ#[1431562](#))

Users no longer get potentially confusing stale data for multipath devices that are not being checked

Previously, when a path device is orphaned (not a member of a multipath device), the device state and checker state displayed with the `show paths` command in the state of the device before it was orphaned. As a result, the `show paths` command showed out of date information on devices that was no longer checking. With this fix, the `show paths` command now displays `undef` as the checker state and `unknown` as the device state of orphaned paths and users no longer get potentially confusing stale data for devices that are not being checked. (BZ#[1402092](#))

The `multipathd` daemon no longer hangs as a result of running the prioritizer on failed paths

Previously, `multipathd` was running the prioritizer on paths that had failed in some cases. Because of this, if `multipathd` was configured with a synchronous prioritizer, it could hang trying to run the prioritizer on a failed path. With this fix, `multipathd` no longer runs the prioritizer when a path has failed and it no longer fails for this reason. (BZ#[1362120](#))

New RAID4 volumes, and existing RAID4 or RAID10 logical volumes after a system upgrade are now correctly activated

After creating RAID4 logical volumes on Red Hat Enterprise Linux version 7.3, or after upgrading a system that has existing RAID4 or RAID10 logical volumes to version 7.3, the system sometimes failed to activate these volumes. With this update, the system activates these volumes successfully. (BZ#[1386184](#))

LVM tools no longer crash due to an incorrect status of PVs

When LVM observes particular types of inconsistencies between the metadata on Physical Volumes (PVs) in a Volume Group (VG), LVM can automatically repair them. Such inconsistencies happen, for example, if a VG is changed while some of its PVs are temporarily invisible to the system and then the PVs reappear.

Prior to this update, when such a repair operation was performed, all the PVs were sometimes temporarily considered to have returned even if this was not the case. As a consequence, LVM tools sometimes terminated unexpectedly with a segmentation fault. With this update, the described problem no longer occurs. (BZ#[1434054](#))

CHAPTER 35. SYSTEM AND SUBSCRIPTION MANAGEMENT

Undercloud no longer fails on a system with no configured repositories

Previously, when the user tried to install the OpenStack **Undercloud** on a system with no configured repositories, the yum package manager required installation of MySQL dependencies which have been already installed. As a consequence, the **Undercloud** install script failed. To fix the bug, yum has been fixed to correctly detect already installed MySQL dependencies. As a result, the **Undercloud** install script no longer fails on a system with no configured repositories. (BZ#1352585)

the yum commands provided by the yum-plugin-verify now set the exit status to 1 if any mismatches are found

The yum commands provided by the **yum-plugin-verify** plug-in returned exit code **0** for any discrepancies found in a package. The bug has been fixed, and the exit status is now set to **1** in case any mismatches are found. (BZ#1406891)

CHAPTER 36. VIRTUALIZATION

SeaBIOS recognizes SCSI devices with a non-zero LUN

Previously, SeaBIOS only recognized SCSI devices when the logical unit number (LUN) was set to zero. Consequently, if a SCSI device was defined with a LUN other than zero, SeaBIOS failed to boot. With this update, SeaBIOS recognizes SCSI devices with LUNs other than zero. As a result, SeaBIOS boots successfully. (BZ#[1020622](#))

The `libguestfs` tools now correctly handle guests where `/usr/` is not on the same partition as root

Previously, the `libguestfs` library did not recognize the guest operating system when the `/usr/` directory was not located on the same partition as the root directory. As a consequence, multiple `libguestfs` tools, such as the `virt-v2v` utility, did not perform as expected when used on such guests. This update ensures that `libguestfs` recognizes guest operating systems when `/usr/` is not on the same partition as root. As a result, the affected `libguestfs` tools perform as expected. (BZ#[1401474](#))

`virt-v2v` can convert Windows guests with corrupted or damaged Windows registries

Previously, the `hivex` library used by `libguestfs` to manipulate the Windows registry could not handle corrupted registries. Consequently, the `virt-v2v` utility was not able to convert Windows guests with corrupted or damaged Windows registries. With this update, `libguestfs` configures `hivex` to be less strict when reading the Windows registry. As a result, `virt-v2v` can now convert most Windows guests with corrupted or damaged Windows registries. (BZ#[1311890](#), BZ#[1423436](#))

Converting Windows guests with non-system dynamic disks using `virt-v2v` now works correctly

Previously, using the `virt-v2v` utility to convert a Windows guest virtual machine with non-system dynamic disks did not work correctly, and the guest were not usable after the conversion. This update fixes the underlying code and thus prevents the described problem.

Note that the conversion of Windows guests using dynamic disks on the system disk (C: drive) is still not supported. (BZ#[1265588](#))

Guests can be converted to Glance images, regardless of the Glance client version

Previously, if the Glance command-line client version 1.0.0 or greater was installed on the `virt-v2v` conversion server, using the `virt-v2v` utility to convert a guest virtual machine to a Glance image failed. With this release, when exporting images, `virt-v2v` directly sets all the properties of images. As a result, the conversion to Glance works regardless of the version of the Glance client installed on the `virt-v2v` conversion server. (BZ#[1374405](#))

Red Hat Enterprise Linux 6.2 - 6.5 guest virtual machines can now be converted using `virt-v2v`

Previously, an error in the SELinux `file_contexts` file in Red Hat Enterprise Linux versions 6.2 - 6.5 prevented conversion of these guests using the `virt-v2v` utility. With this update, `virt-v2v` automatically fixes the error in the SELinux `file_contexts` file. As a result, Red Hat Enterprise Linux 6.2-6.5 guest virtual machines can now be converted using `virt-v2v`. (BZ#[1374232](#))

Btrfs entries in `/etc/fstab` are now parsed correctly by `libguestfs`

Previously, Btrfs sub-volume entries with more than one comma-separated option in `/etc/fstab` were not parsed properly by `libguestfs`. Consequently, Linux guest virtual machines with these configurations could not be inspected, and the `virt-v2v` utility could not convert them. With this

update, `libguestfs` parses Btrfs sub-volume entries with more than one comma-separated option in `/etc/fstab` correctly. As a result, these entries can be inspected and converted by `virt-v2v`. (BZ#1383517)

libguestfs can now correctly open libvirt domain disks that require authentication

Previously, when adding disks from a `libvirt` domain, `libguestfs` did not read any disk secrets. Consequently, `libguestfs` could not open disks that required authentication. With this update, `libguestfs` reads secrets of disks in `libvirt` domains, if present. As a result, `libguestfs` can now correctly open disks of `libvirt` domains that require authentication. (BZ# 1392798)

Converted Windows UEFI guests boot properly

Previously, when converting Windows 8 UEFI guests, virtio drivers were not installed correctly. Consequently, the converted guests did not boot. With this update, virtio drivers are installed correctly in Windows UEFI guests. As a result, converted Windows UEFI guests boot properly. (BZ#1431579)

The virt-v2v utility now ignores proxy environment variables consistently

Prior to this update, when using the `virt-v2v` utility to convert a VMware guest virtual machine, `virt-v2v` used the proxy environment variables for some connections to VMware, but not for others. This in some cases caused conversions to fail. Now, `virt-v2v` ignores all proxy environment settings during the conversion, which prevents the described problem. (BZ#1354507)

virt-v2v only copies rhev-apt.exe and rhsvany.exe when needed

Previously, `virt-v2v` always copied the `rhev-apt.exe` and `rhsvany.exe` files when converting Windows guests. Consequently, they were present in the converted Windows guests, even when they were not needed. With this update, `virt-v2v` only copies these files when they are needed in the Windows guest. (BZ#1161019)

Guests with VLAN over a bonded interface no longer stop passing traffic after a failover

Previously, on guest virtual machines with VLAN configured over a bonded interface that used `ixgbe` virtual functions (VFs), the bonded network interface stopped passing traffic when a failover occurred. The hypervisor console also logged this error as a `requested MACVLAN filter but is administratively denied` message. With this update ensures that failovers are handled correctly and thus prevents the described problem. (BZ#1379787)

virt-v2v imports OVAs that do not have the <ovf:Name> attribute

Previously, the `virt-v2v` utility rejected the import of Open Virtual Appliances (OVAs) without the `<ovf:Name>` attribute. As a consequence, the `virt-v2v` utility did not import OVAs exported by Amazon Web Services (AWS). In this release, if the `<ovf:Name>` attribute is missing, `virt-v2v` uses the base name of the disk image file as the name of the virtual machine. As a result, the `virt-v2v` utility now imports OVAs exported by AWS. (BZ#1402301)

PART III. TECHNOLOGY PREVIEWS

This part provides a list of all Technology Previews available in Red Hat Enterprise Linux 7.4.

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

CHAPTER 37. 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 38. 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 39. 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](#))

CHAPTER 40. COMPILER AND TOOLS

Shenandoah garbage collector

The new, low pause time Shenandoah garbage collector, is now available as a Technology Preview for OpenJDK 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#1400306)

CHAPTER 41. FILE SYSTEMS

The CephFS kernel client is now available

Starting with Red Hat Enterprise Linux 7.3, the Ceph File System (CephFS) kernel module enables, as a Technology Preview, 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. For more information, see the Ceph File System Guide for Red Hat Ceph Storage 2: <https://access.redhat.com/documentation/en/red-hat-ceph-storage/2/single/ceph-file-system-guide-technology-preview> (BZ#1205497)

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 and Block Layout Support

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

In addition, Red Hat Enterprise Linux 7.4 includes the Technology Preview of the pNFS SCSI layout. This feature is similar to pNFS block layout, but limited only to SCSI devices, so it is easier to use. Therefore, Red Hat recommends the use of pNFS SCSI layout rather than the pNFS block layout. (BZ#1111712)

OverlayFS

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

OverlayFS remains a Technology Preview in Red Hat Enterprise Linux 7.4 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 7.3 release. For details, see **Non-standard behavior** in the **Documentation/filesystems/overlayfs.txt** file. (BZ#1206277)

pNFS SCSI layouts client and server support is now provided

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)

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 introduces 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)

CHAPTER 42. HARDWARE ENABLEMENT

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

Two new packages have been added to Red Hat Enterprise Linux to support the Trusted Computing Group's Trusted Platform Module (TPM) 2.0 hardware as a Technology Preview:

- 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#[1275027](#), BZ#1275029)

New package: `tss2`

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

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)

CHAPTER 43. INSTALLATION AND BOOTING

Multi-threaded xz compression in rpm-build

Compression can take long time for highly parallel builds as it currently uses only one core. This is problematic especially for continuous integration of large projects that are built on hardware with many cores.

This feature, which is provided as a Technology Preview, enables 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`. This can be done by editing the `/usr/lib/rpm/macros` file or by declaring the macro within the spec file or at the command line. (BZ#1278924)

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 2.12

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.

With Red Hat Enterprise Linux 7.4, the `criu` packages have been upgraded to upstream version 2.12, which provides a number of bug fixes and enhancements over the previous version. (BZ#1400230)

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)

KASLR as a Technology Preview

Kernel address space layout randomization (KASLR) is now available as a Technology Preview. 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.

Note that KASLR code is now compiled in the Linux kernel, but it is disabled by default. If you want to use it, add the `kaslr` kernel option to the kernel command line to enable it explicitly. (BZ#1449762)

Updated NFSv4 pNFS clients with flexible file layout

Flexible file layout on NFSv4 clients was first introduced in Red Hat Enterprise Linux 7.2 as a Technology Preview. Red Hat Enterprise Linux 7.4 adds updates to this feature, however, it is still being offered as a Technology Preview.

NFSv4 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 NFS flexible file layout. (BZ#1349668)

CUIR enhanced scope detection

The Linux support for Control Unit Initiated Reconfiguration (CUIR) enables concurrent storage service with no or minimized down time. 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 as a Technology Preview. (BZ#1274456)

SCSI-MQ as a Technology Preview in the qla2xxx 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)

CHAPTER 45. REAL-TIME KERNEL

New scheduler class: SCHED_DEADLINE

This update introduces the **SCHED_DEADLINE** scheduler class for the real-time kernel as a Technology Preview. The new 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

The `nftables` and `libnftnl` packages are available as a Technology Preview since Red Hat Enterprise Linux 7.3.

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

New packages: `ansible`

Red Hat Enterprise Linux System Roles, now 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.

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

Notes:

- Currently, `Ansible` is not a part of the Red Hat Enterprise Linux FIPS validation process. We hope to address this in future releases.
- `Ansible` is being included as an unsupported runtime dependency. (BZ#1313263)

CHAPTER 48. SECURITY

The tang-nagios and clevis-udisk2 subpackages available as a Technology Preview

The tang and clevis packages, which are part of the Red Hat Enterprise Linux Network Bound Disk Encryption (NBDE) project, contain also the tang-nagios and clevis-udisk2 subpackages. These subpackages are provided only as a Technology Preview. (BZ#1467338)

USBGuard is now available for IBM Power as a Technology Preview

The usbguard packages, which provide system protection against intrusive USB devices, are now available. With this update, the **USBGuard** software framework for the IBM Power architectures is provided as a Technology Preview. Full support is targeted for a later release of Red Hat Enterprise Linux.

Note that USB is not supported on IBM z Systems, and the **USBGuard** framework cannot be provided on those systems. (BZ#[1467369](#))

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 7.4. (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 ixgbevf 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)

The `ibmvnic` Device Driver has been added

The `ibmvnic` Device Driver was introduced as a Technology Preview in Red Hat Enterprise Linux 7.3 for IBM POWER architectures. vNIC (Virtual Network Interface Controller) is a new 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#947163)

`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. (BZ#1441197)

`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)

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 7.4.

CHAPTER 51. NEW DRIVERS

Storage Drivers

- nvme-fabrics
- nvme-rdma
- nvmet
- nvmet-rdma
- nvme-loop
- qedi
- qedf

Network Drivers

- qedr
- rdma_rxe
- ntb_transport
- ntb_perf
- mdev
- vfio_mdev
- amd-xgbe
- atlantic
- libcxgb
- ena
- rocker
- amd8111e
- nfp
- mlxsw_core
- mlxsw_i2c
- mlxsw_spectrum
- mlxsw_pci
- mlxsw_switchx2
- mlxsw_switchib

- mlxsw_minimal

Graphics Drivers and Miscellaneous Drivers

- ccp
- chcr
- uio_hv_generic
- usbip-core
- vhost_vsock
- tpm_tis_spi
- gpio-amdpt
- joydev
- sdio_uart
- ptp_kvm
- mei_wdt
- dell-rbtn
- dell-smo8800
- intel-hid
- dell-smbios
- skx_edac
- kvmgt
- pinctrl-intel
- pinctrl-sunrisepoint
- pinctrl-amd
- dax_pmem
- dax
- nfit
- ledtrig-usbport

CHAPTER 52. UPDATED DRIVERS

Storage Driver Updates

- The aacraid driver has been updated to version 1.2.1[50792]-custom.
- The lpfc driver has been updated to version 0:11.2.0.6.
- The vmw_pvscsi driver has been updated to version 1.0.7.0-k.
- The megaraid_sas driver has been updated to version 07.701.17.00-rh1.
- The bfa driver has been updated to version 3.2.25.1.
- The hpsa driver has been updated to version 3.4.18-0-RH1.
- The be2iscsi driver has been updated to version 11.2.1.0.
- The qla2xxx driver has been updated to version 8.07.00.38.07.4-k1.
- The mpt2sas driver has been updated to version 20.103.00.00.
- The mpt3sas driver has been updated to version 15.100.00.00.

Network Driver Updates

- The ntb driver has been updated to version 1.0.
- The igbvf driver has been updated to version 2.4.0-k.
- The igb driver has been updated to version 5.4.0-k.
- The ixgbev driver has been updated to version 3.2.2-k-rh7.4.
- The i40e driver has been updated to version 1.6.27-k.
- The fm10k driver has been updated to version 0.21.2-k.
- The i40evf driver has been updated to version 1.6.27-k.
- The ixgbe driver has been updated to version 4.4.0-k-rh7.4.
- The be2net driver has been updated to version 11.1.0.0r.
- The qede driver has been updated to version 8.10.10.21.
- The qlge driver has been updated to version 1.00.00.35.
- The qed driver has been updated to version 8.10.10.21.
- The bna driver has been updated to version 3.2.25.1r.
- The bnxt driver has been updated to version 1.7.0.
- The enic driver has been updated to version 2.3.0.31.
- The fjes driver has been updated to version 1.2.

- The hpwdt driver has been updated to version 1.4.02.

Graphics Driver and Miscellaneous Driver Updates

- The vmwgfx driver has been updated to version 2.12.0.0.
- The hpilo driver has been updated to version 1.5.0.

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 7.4.

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.

Deprecated packages related to Identity Management

The following packages are 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]
[a] System Security Services Daemon (SSSD) contains enhanced smart card functionality. [b] For details on migrating from pam_krb5 to sssd, see Migrating from pam_krb5 to sssd in the upstream SSSD documentation. [c] Red Hat Directory Server requires a valid Directory Server subscription.	

Deprecated Insecure Algorithms and Protocols

Algorithms that provide cryptographic hashes and encryption as well as cryptographic protocols have a lifetime after which they are considered either too risky to use or plain insecure. See the [Enhancing the Security of the Operating System with Cryptography Changes in Red Hat Enterprise Linux 7.4](#) Knowledgebase article on the Red Hat Customer Portal for more information.

Weak ciphers and algorithms are no longer used by default in OpenSSH

With this update, the **OpenSSH** library removes several weak ciphers and algorithms from default configurations. However, backward compatibility is ensured in most cases.

The following have been removed from the **OpenSSH** server and client:

- Host key algorithms:
 - ssh-rsa-cert-v00@openssh.com
 - ssh-dss-cert-v00@openssh.com
- Ciphers:
 - arcfour256
 - arcfour128
 - arcfour
 - rijndael-cbc@lysator.liu.se
- MACs:
 - hmac-md5
 - hmac-md5-96
 - hmac-md5-96-etm@openssh.com
 - hmac-md5-etm@openssh.com
 - hmac-ripemd160
 - hmac-ripemd160-etm@openssh.com
 - hmac-ripemd160@openssh.com
 - hmac-sha1-96
 - hmac-sha1-96-etm@openssh.com

The following have been removed from the **OpenSSH** client:

- Ciphers:
 - blowfish-cbc
 - cast128-cbc
 - 3des-cbc

OpenSSH no longer uses the SHA-1-based key exchange algorithms in FIPS mode

This update removes the SHA-1-based key exchange algorithms from the default list in FIPS mode. To enable those algorithms, use the following configuration snippet for the `~/.ssh/config` and `/etc/ssh/sshd_config` files:

```
KexAlgorithms=+diffie-hellman-group14-sha1,diffie-hellman-group-exchange-sha1
```

The SSH-1 protocol has been removed from the OpenSSH server

SSH-1 protocol support has been removed from the **OpenSSH** server. For more information, see the [The server-side SSH-1 protocol removal from RHEL 7.4](#) Knowledgebase article.

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

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

Additionally, the default algorithm for generating digital signatures has been changed from SHA-1 to SHA-256. The verification of SHA-1 signatures is still enabled for legacy purposes.

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

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

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

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

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

This update prevents **OpenSSL** clients from connecting to servers with Diffie-Hellman (DH) parameters shorter than 1024 bits. This ensures that clients using **OpenSSL** are not susceptible to vulnerabilities, such as Logjam.

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

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

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

SSL 2.0 support has been completely removed from **OpenSSL**

The SSL protocol version 2.0, which is considered insecure for more than seven years, was deprecated by RFC 6176 in 2011. In Red Hat Enterprise Linux, support of SSL 2.0 was already disabled by default. With this update, SSL 2.0 support has been removed completely. The **OpenSSL** library API calls that use this protocol version now return an error message.

EXPORT cipher suites in **OpenSSL have been deprecated**

This change removes support for EXPORT cipher suites from the **OpenSSL** toolkit. Disabling these weak cipher suites ensures that clients using **OpenSSL** are not susceptible to vulnerabilities, such as FREAK. EXPORT cipher suites are no longer required in any TLS protocol configurations.

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

This change prevents GNU Transport Layer Security (GnuTLS) clients from connecting to servers with Diffie-Hellman (DH) parameters shorter than 1024 bits. This ensures that clients using **GnuTLS** are not susceptible to vulnerabilities, such as Logjam.

In applications that accept a priority string from the user or configuration directly, this change can be reverted by appending the priority string `%PROFILE_VERY_WEAK` to the used priority string.

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

This change prevents Network Security Services (NSS) clients from connecting to servers with Diffie-Hellman (DH) parameters shorter than 1024 bits. This ensures that clients using NSS are not susceptible to vulnerabilities, such as Logjam.

The system administrator can enable shorter DH parameter support by modifying the `/etc/pki/nss-legacy/nss-rhel7.config` policy configuration file to:

```
library=
name=Policy
NSS=flags=policyOnly,moduleDB
config="allow=DH-MIN=767:DSA-MIN=767:RSA-MIN=767"
```

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

EXPORT cipher suites in NSS have been deprecated

This change removes support for EXPORT cipher suites in the Network Security Services (NSS) library. Disabling these weak cipher suites protects against vulnerabilities, such as FREAK. EXPORT cipher suites are not required in any TLS protocol configuration.

Legacy CA certificates removed from the ca-certificates package

Previously, to allow older versions of the `GnuTLS`, `OpenSSL`, and `glib-networking` libraries to remain compatible with the Public Key Infrastructure (PKI), the `ca-certificates` package included a set of legacy CA certificates with 1024-bit RSA keys as trusted by default.

Since Red Hat Enterprise Linux 7.4, updated versions of `OpenSSL`, `GnuTLS`, and `glib-networking` are available, which are able to correctly identify a replacement of root CA certificates. Trusting these legacy CA certificates is no longer required for public web PKI compatibility.

The legacy configuration mechanism, which could previously be used to disable the legacy CA certificates, is no longer supported; the list of legacy CA certificates has been changed to empty.

The `ca-legacy` tool is still available and it also keeps current configuration settings for potential future reuse.

coolkey replaced with opensc

The `OpenSC` library implements the `PKCS#11` API and replaces the `coolkey` packages. In Red Hat Enterprise Linux 7, the CoolKey Applet functionality is also provided by the `opensc` package.

The `coolkey` package will remain supported for the lifetime of Red Hat Enterprise Linux 7, but new hardware enablement will be provided through the `opensc` package.

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, which provides a library and a small daemon program that can monitor and filter incoming requests for **systat**, **finger**, **FTP**, **telnet**, **rlogin**, **rsh**, **exec**, **tftp**, **talk**, **sshd**, and other network services, has been deprecated.

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.wrap_socket()** function instead.

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 `ipson` 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 `ipson` 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; see [Chapter 8, Desktop](#) for more information.

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

- `3w-9xxx`
- `3w-sas`
- `mptbase`
- `mptctl`
- `mptsas`
- `mptscsih`
- `mptspi`

- `mvsas`
- `qla3xxx`
- The following controllers from the `megaraid_sas` driver have been deprecated:
 - Dell PERC5, PCI ID 0x15
 - SAS1078R, PCI ID 0x60
 - SAS1078DE, PCI ID 0x7C
 - SAS1064R, PCI ID 0x411
 - VERDE_ZCR, PCI ID 0x413
 - SAS1078GEN2, PCI ID 0x78
- The following adapters from the `qla2xxx` driver have been deprecated:
 - ISP24xx, PCI ID 0x2422
 - ISP24xx, PCI ID 0x2432
 - ISP2422, PCI ID 0x5422
 - QLE220, PCI ID 0x5432
 - QLE81xx, PCI ID 0x8001
 - QLE10000, PCI ID 0xF000
 - QLE84xx, PCI ID 0x8044
 - QLE8000, PCI ID 0x8432
 - QLE82xx, PCI ID 0x8021
- The following Ethernet adapter controlled by the `be2net` driver has been deprecated:
 - TIGERSHARK NIC, PCI ID 0x0700
- The following controllers 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 Emulex boards 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 controllers from the mentioned drivers that are not listed here remain unchanged.

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 portion of 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 userspace 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 7.4.

CHAPTER 54. AUTHENTICATION AND INTEROPERABILITY

sudo unexpectedly denies access when performing group lookups

This problem occurs on systems that meet all of these conditions:

- A group name is configured in a `sudoers` rule available through multiple Name Service Switch (NSS) sources, such as `files` or `sss`.
- The NSS priority is set to local group definitions. This is true when the `/etc/nsswitch.conf` file includes the following line:

```
sudoers: files sss
```

- The `sudo` Defaults option named `match_group_by_gid` is set to `true`. This is the default value for the option.

Because of the NSS source priority, when the `sudo` utility tries to look up the GID of the specified group, `sudo` receives a result that describes only the local group definition. Therefore, if the user is a member of the remote group, but not the local group, the `sudoers` rule does not match, and `sudo` denies access.

To work around this problem, choose one of the following:

- Explicitly disable the `match_group_by_gid` Defaults for `sudoers`. Open the `/etc/sudoers` file, and add this line:

```
Defaults !match_group_by_gid
```

- Configure NSS to prioritize the `sss` NSS source over `files`. Open the `/etc/nsswitch.conf` file, and make sure it lists `sss` before `files`:

```
sudoers: sss files
```

This ensures that `sudo` permits access to users that belong to the remote group. (BZ#1293306)

The KCM credential cache is not suitable for a large number of credentials in a single credential cache

If the credential cache contains too many credentials, Kerberos operations, such as `klist`, fail due to a hardcoded limit on the buffer used to transfer data between the `sssd-kcm` component and the `sssd-secrets` component.

To work around this problem, add the `ccache_storage = memory` option in the `[kcm]` section of the `/etc/sss/sss.conf` file. This instructs the `kcm` responder to only store the credential caches in-memory, not persistently. Note that if you do this, restarting the system or `sssd-kcm` clears the credential caches. (BZ#1448094)

The sssd-secrets component crashes when it is under load

When the `sssd-secrets` component receives many requests, the situation triggers a bug in the Network Security Services (NSS) library that causes `sssd-secrets` to terminate unexpectedly. However, the `systemd` service restarts `sssd-secrets` for the next request, which means that the denial of service is only temporary. (BZ#1460689)

SSSD does not correctly handle multiple certificate matching rules with the same priority

If a given certificate matches multiple certificate matching rules with the same priority, the System Security Services Daemon (SSSD) uses only one of the rules. As a workaround, use a single certificate matching rule whose LDAP filter consists of the filters of the individual rules concatenated with the | (or) operator. For examples of certificate matching rules, see the `sss-certmap(5)` man page. (BZ#1447945)

SSSD can look up only unique certificates in ID overrides

When multiple ID overrides contain the same certificate, the System Security Services Daemon (SSSD) is unable to resolve queries for the users that match the certificate. An attempt to look up these users does not return any user. Note that looking up users by using their user name or UID works as expected. (BZ#1446101)

The `ipa-adviser` command does not fully configure smart card authentication

The `ipa-adviser config-server-for-smart-card-auth` and `ipa-adviser config-client-for-smart-card-auth` commands do not fully configure the Identity Management (IdM) server and client for smart card authentication. As a consequence, after running the script that the `ipa-adviser` command generated, smart card authentication fails. To work around the problem, see the manual steps for the individual use case in the Linux Domain Identity, Authentication, and Policy Guide: https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/Linux_Domain_Identity_Authentication_and_Policy_Guide/smart_cards.html (BZ#1455946)

The `libwbclient` library fails to connect to Samba shares hosted on Red Hat Enterprise Linux 7.4

The interface between Samba and the System Security Services Daemon's (SSSD) Winbind plug-in implementation changed. However, this change is missing in SSSD. As a consequence, systems that use the SSSD `libwbclient` library instead of the Winbind daemon fail to access shares provided by Samba running on Red Hat Enterprise Linux 7.4. There is no workaround available, and Red Hat recommends to not upgrade to Red Hat Enterprise 7.4 if you are using the `libwbclient` library without running the Winbind daemon. (BZ#1462769)

Certificate System subsystems experience communication problems with TLS_ECDHE_RSA_* ciphers and certain HSMs

When certain HSMs are used while TLS_ECDHE_RSA_* ciphers are enabled, subsystems experience communication problems. The issue occurs in the following scenarios:

- When a CA has been installed and a second subsystem is being installed and tries to contact the CA as a security domain, thus preventing the installation from succeeding.
- While performing a certificate enrollment on the CA, when archival is required, the CA encounters the same communication problem with the KRA. This scenario can only occur if the offending ciphers were temporarily disabled for the installation.

To work around this problem, keep the TLS_ECDHE_RSA_* ciphers turned off if possible. Note that while the Perfect Forward Secrecy provides added security by using the TLS_ECDHE_RSA_* ciphers, each SSL session takes about three times longer to establish. Also, the default TLS_RSA_* ciphers are adequate for the Certificate System operations. (BZ#1256901)

CHAPTER 55. COMPILER AND TOOLS

Performance of regular expressions cannot be boosted with the JIT technique if executable stack is disabled

When the **SELinux** policy disallows executable stack, the **PCRE** library cannot use JIT compilation to speed up regular expressions. As a result, attempting JIT compilation for regular expressions is ignored and their performance is not boosted.

To work around this problem, amend the **SELinux** policy with a rule for enabling the **execmem** action on affected **SELinux** domains to enable JIT compilation. Some of the rules are already provided and can be enabled by specific **SELinux** booleans. To list these booleans, see the output of the following command:

```
getsebool -a | grep execmem
```

An alternative workaround is changing application code to not request JIT compilation with calls to the `pcre_study()` function. (BZ#[1290432](#))

Memory leaks occur when certain applications fail to exit after unloading the Gluster libraries

Gluster consists of many internal components and different translators that implement functions and features. The **gfsapi** access method was added to integrate Gluster tightly with applications. However, not all components and translators are designed to be unloaded in running applications. As a consequence, programs that do not exit after unloading the Gluster libraries are unable to release some of the memory allocations that are performed internally by Gluster.

To reduce the amount of memory leaks, prevent applications from calling the `gfs_init()` and `gfs_fini()` functions whenever possible. To release the leaked memory, you must restart long-running applications. (BZ#[1409773](#))

URL to DISA SRGs is incorrect

The SCAP Security Guide (SSG) rules refer to Defense Information Systems Agency Security Requirement Guides (DISA SRGs). Connecting to the URL fails with an **404 - not found** error. As a consequence, users have no direct reference to SRGs. To work around the problem, use the new URL: <http://iase.disa.mil/stig/os/general/Pages/index.aspx/> (BZ#[1464899](#))

The `ensure_gpgcheck_repo_metadata` rule fails

During remediation of the `ensure_gpgcheck_repo_metadata` rule, certain profiles update the `yum.conf` file to enable the `repo_gpgcheck` option. Red Hat does not currently provide signed repository metadata. As a consequence, the `yum` utility is no longer able to install any package from official repositories. To work around the problem, use a tailoring file to remove `ensure_gpgcheck_repo_metadata` from the profile. If remediation already breaks the system, update `yum.conf` and set `repo_gpgcheck` to `0`. (BZ#[1465677](#))

The SSG `pam_faillock` module utilization check incorrectly accepts `default=die`

The SCAP Security Guide (SSG) `pam_faillock` module utilization check incorrectly accepts the `default=die` option. Consequently, when a user authentication using the `pam_unix` module fails, the `pam` stack evaluation stops immediately without incrementing the counter of `pam_faillock`. To work around this problem, do not use `default=die` before the `authfail` option. This ensures that the `pam_faillock` counter is incremented properly. (BZ#[1448952](#))

CHAPTER 56. DESKTOP

Updating totem alone fails

Explicit dependency between the totem and gstreamer1-plugins-bad-free packages is missing. Consequently, when trying to update just the totem package, the operation fails. To work around this problem, do not update the totem package by itself and rely on the system update instead. (BZ#[1451211](#))

The operating system always assumes Wacom Expresskeys Remote mode 1 when booting

The Wacom Expresskeys Remote (EKR) is a standalone device and thus can be switched to any operating mode when the operating system (OS) boots. Currently, however, the OS always assumes that the EKR is set to mode 1 upon boot. Consequently, the EKR is not synchronized with the OS when the EKR mode was not set to 1 prior to system boot. To work around this problem, set the EKR to mode 1 before starting the OS. (BZ#[1458351](#))

Cannot install downloaded RPM files from Nautilus

When an RPM file is double clicked in the **Nautilus** file manger, the following error is returned instead of the file being installed:

```
Sorry, this did not work, File is not supported
```

This happens because the **yum** backend to **PackageKit** does not support getting details about local files.

This problem can be worked around by either installing **gnome-packagekit** to handle the double-click action or by manually installing the files using **yum**. (BZ#[1434477](#))

Yelp does not correctly display HTML formatted files

In earlier versions of **yelp**, HTML formatted files could be displayed. With version 3.22, this functionality no longer works, returning an **Unknown Error** with the qualifying text **URL cannot be shown**.

There is no workaround at this time, as the problem may be related to architectural changes in **yelp** itself.

System administrators should note that **yelp** does not support this use case and expects **Mallard** or **Docbook** data as input.

Alternative methods for displaying HTML formatted content should be considered. (BZ#[1443179](#))

Automatic modesetting fails when attaching monitors with some AMD hardware

With some configurations, adding an additional monitor to a system using AMD hardware can fail to activate the new hardware automatically.

The problem is currently under investigation.

To workaround the issue, system administrators should manually call **xrandr(1)** to enable the monitor. (BZ#[1393951](#))

Gnome Documents can not display some documents when installed without LibreOffice due to a missing dependency

Gnome Documents uses libraries provided by the **LibreOffice** suite for rendering certain types of documents, such as OpenDocument Text or Open Office XML formats. However, the required libraries

(libreoffice-filters) are missing from the dependency list of the `gnome-documents` package. Therefore, if you install **Gnome Documents** on a system that does not have **LibreOffice**, the aforementioned document types can not be rendered.

To work around this problem, install `libreoffice-filters` package manually, even if you do not plan to use LibreOffice itself. (BZ#[1466164](#))

Application Installer displays packages even though they can not be installed on big endian architectures

When you use the **Application Installer** graphical package installer (the `gnome-software` package) on a big-endian system such as IBM Power Systems or IBM z Systems, some of the packages listed as available will not be possible to install, and an attempt to do so results in an **installing not available** error message. This is a known issue caused by package metadata currently being generated only for 64-bit AMD and Intel-compatible (little-endian) systems, and assuming all packages are also available on big-endian architectures, which is not the case.

There is no workaround to this problem; however, the error message has no consequences other than the package not being installable. (BZ#[1464139](#))

The Add/Remove Software tool (gpk-application) does not use a newly imported key on the first try

When using the **Add/Remove Software** graphical interface in **GNOME** to install a package signed by a key that was not already imported, the tool displays a prompt allowing you to import the key. However, even if the key is imported, the installation is still unsuccessful due to a bug that prevents using the key immediately. To work around the problem, install the same package again; at that point the key is already imported from the previous attempt, and the installation succeeds. (BZ#[1387181](#))

Resizing a display of a virtual machine with multiple displays using multiple PCI devices causes X to crash

A bug in the `QXL` driver (`xorg-x11-drv-qxl`) causes the `X.Org` display server on a virtual machine to crash upon display resize if the virtual machine has multiple displays configured to use multiple PCI devices. Make sure guest virtual machines running Red Hat Enterprise Linux with multiple monitors are configured to use a single PCI device. In Red Hat Virtualization, this setting is controlled by the **Single PCI Device** check box under **Edit -> Console**, and is enabled by default. (BZ#[1428340](#))

Nautilus does not hide icons in the GNOME Classic Session

The **GNOME Tweak Tool** to show or hide icons in the `gnome` session, where the icons are hidden by default, is ignored in the **GNOME Classic Session**. As a result, it is not possible to hide icons in the **GNOME Classic Session** even though the **GNOME Tweak Tool** displays this option. (BZ#[1474852](#))

Incorrect dependency in flatpak

Due to a wrong dependency in the `flatpak` package, the user can encounter the following error:

```
flatpak: error while loading shared libraries: libostree-1.so.1: cannot
open shared object file: No such file or directory
```

To work around this problem, install the `flatpak-libs` package. Alternatively, instead of initially installing just `flatpak`, install both of the packages by running:

```
sudo yum -y install flatpak flatpak-libs
```

(BZ#[1476905](#))

Firefox does not start after update

After the upgrade to `firefox-52.1.2-el7.x86_64` and further, the browser does not start in some cases. This is due to the `nspr` and `nss` packages not being updated from the Red Hat Enterprise Linux 7.4 batch. To work around this problem, update the `nspr` and `nss` packages from Red Hat Enterprise 7.4 release. Another possible workaround is to downgrade Firefox, but this option is not recommended. As a result, it is possible to start the Firefox web browser again. (BZ#[1455798](#))

Limited support for visuals in Xorg

In the Xorg server, only `TrueColor` and `DirectColor` visuals at depth 16 or higher are supported for hardware drivers. Legacy applications that need a `PseudoColor` visual can be run against the Xephyr nested X server, which implements `PseudoColor` translation when displayed on a `TrueColor` screen. (BZ#1185690)

CHAPTER 57. FILE SYSTEMS

NetApp storage appliances serving NFSv4 are advised to check their configuration

Note that features can be enabled or disabled on a per-minor version basis when using NetApp storage appliances that serve NFSv4.

It is recommended to verify the configuration to ensure that the appropriate features are enabled as desired, for example by using the following Data ONTAP command:

```
vserver nfs show -vserver <vserver-name> -fields v4.0-acl,v4.0-read-  
delegation,v4.0-write-delegation,v4.0-referrals,v4.0-migration,v4.1-  
referrals,v4.1-migration,v4.1-acl,v4.1-read-delegation,v4.1-write-  
delegation
```

(BZ#[1450447](#))

CHAPTER 58. HARDWARE ENABLEMENT

The i40e driver rejects the most general HWTSTAMP filter

Because of disabling L4 timestamping (UDP) in the security fix for Intel Ethernet Controller X710 and XL710 family described in the INTEL-SA-00063 advisory, the i40e device driver rejects the most general HWTSTAMP filter. This problem only affects the Intel X710 devices, not the newer X722 device. (BZ#1431964)

CHAPTER 59. INSTALLATION AND BOOTING

FIPS mode unsupported when installing from an HTTPS kickstart source

Installation images do not support FIPS mode during installation with an HTTPS kickstart source. As a consequence, it is currently impossible to install a system with the `fips=1` and `inst.ks=https://<location>/ks.cfg` options added to the command line. (BZ# [1341280](#))

PXE boot with UEFI and IPv6 displays the GRUB2 shell instead of the operating system selection menu

When the **Pre-Boot Execution Environment (PXE)** starts on a client configured with UEFI and IPv6, the boot menu configured in the `/boot/grub/grub.cfg` file is not displayed. After a timeout, the **GRUB2** shell is displayed instead of the configured operating system selection menu. (BZ#1154226)

Specifying a driverdisk partition with non-alphanumeric characters generates an invalid output Kickstart file

When installing Red Hat Enterprise Linux using the **Anaconda** installer, you can add a driver disk by including a path to the partition containing the driver disk in the Kickstart file. At present, if you specify the partition by LABEL or CDLABEL which has non-alphanumeric characters in it, for example:

```
driverdisk "CDLABEL=Fedora 23 x86_64:/path/to/rpm"
```

the output Kickstart file created during the **Anaconda** installation will contain incorrect information. To work around this problem, use only alphanumeric characters when specifying the partition by LABEL or CDLABEL. (BZ#[1452770](#))

The Scientific Computing variant is missing packages required for certain security profiles

When installing the **Red Hat Enterprise Linux for Scientific Computing** variant, also known as **Compute Node**, you can select a security profile similarly to any other variant's installation process. However, since this variant is meant to be minimal, it is missing packages which are required by certain profiles, such as **United States Government Configuration Baseline**. If you select this profile, the installer displays a warning that some packages are missing.

The warning allows you to continue the installation despite missing packages, which can be used to work around the problem. The installation will complete normally, however, note that if you install the system despite the warning, and then attempt to run a security scan after the installation, the scan will report failing rules due to these missing packages. This behavior is expected. (BZ#1462647)

CHAPTER 60. KERNEL

kexec fails when secondary cores do not offline

Under certain circumstances, secondary-core offlining fails on AppliedMicro X-Gene platforms like HP ProLiant m400 and AppliedMicro Mustang. As a consequence, the kernel sometimes fails to trigger the `kdump` crash dump mechanism through `kexec` when a kernel panic occurs. As a result, the kernel crash dump file is not saved. (BZ#1218374)

File-system corruption due to incorrect flushing of cache has been fixed but I/O operations can be slower

Due to a bug in the `megaraid_sas` driver, file-system corruption previously occurred in some cases when the file system was used with a disk-write back cache during system shutdown, reboot, or power loss. This update fixes `megaraid_sas` to transfer the flush cache commands correctly to the raid card. As a result, if you also update the raid card firmware, the file-system corruption no longer occurs under the described circumstances.

With Broadcom `megaraid_sas` raid adapter, you can check the functionality in the system log (`dmesg`). The proper functionality is indicated by the following text string:

```
FW supports sync cache Yes
```

Note that this fix can slow down I/O operations because the cache is now flushed properly. (BZ#1380447)

Wacom Cintiq 12WX is not redetected when unplugged and plugged in quickly

When unplugging and quickly plugging in Wacom Cintiq 12WX within the same USB port, the tablet is currently not recognized. To work around this problem, wait 3-5 seconds before plugging the tablet back in. (BZ#1458354)

Installing to some IBM POWER8 machines using a Virtual DVD fails when starting GUI

Red Hat Enterprise Linux 7.4 can fail to install on some IBM POWER8 hardware (including S822LC machines) while starting the Anaconda GUI.

The problem is characterized by errors starting X11, followed by a `Pane is dead` message in the Anaconda screen.

The workaround is to append `inst . text` to the kernel command line and install in text mode.

This issue is confined to Virtual DVD installations, additional testing with the netboot image allows GUI installation. (BZ#1377857)

Entering full screen mode using a keyboard shortcut causes display problems on VMWare ESXi 5.5

When using Red Hat Enterprise Linux 7.4 as a virtual machine guest running on a VMWare ESXi 5.5 host, pressing `Ctrl+Alt+Enter` to enter full screen mode in the console causes the display to become unusable. At the same time, errors such as the following example are saved into the system log (`dmesg`):

```
[drm:vmw_cmdbuf_work_func [vmwgfx]] *ERROR* Command buffer error.
```

To work around this problem, shut down the virtual machine, open its `.vmx` configuration file, and add or modify the following parameters:

■

```
svga.maxWidth = X  
svga.maxHeight = Y  
svga.vramSize = "X * Y * 4"
```

In the above, replace X and Y with the horizontal and vertical resolution of your screen. The `svga.vramSize` parameter takes a value that is equal to X times Y times 4. An example setup for a screen with a resolution of 1920x1080 therefore is:

```
svga.maxWidth = 1920  
svga.maxHeight = 1080  
svga.vramSize = "8294400"
```

Note that VMWare ESXi 5.5 is the only version reported to encounter this bug; other versions can enter full screen mode without problems. (BZ#1451242)

KSC currently does not support xz compression

The Kernel module Source Checker (the `ksc` tool) is currently unable to process the `xz` compression method and reports the following error:

```
Invalid architecture, supported architectures are x86_64, ppc64, s390x
```

Until this limitation is resolved, system administrators should manually uncompress any third party modules using `xz` compression, before running the `ksc` tool. (BZ#1463600)

CHAPTER 61. 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 62. SECURITY

certutil does not return the NSS database password requirements in FIPS mode

When creating a new Network Security Services (NSS) database with the `certutil` tool, the user has nowhere to find out what the database password requirements are when running in FIPS mode. The prompt message does not provide password requirements, and `certutil` returns only a generic error message:

```
certutil: could not authenticate to token NSS FIPS 140-2 Certificate DB.:  
SEC_ERROR_IO: An I/O error occurred during security authorization.
```

(BZ#[1401809](#))

systemd-importd runs as init_t

The `systemd-importd` service is using the `NoNewPrivileges` security flag in the `systemd` unit file. This blocks the SELinux domain transition from the `init_t` to `systemd_importd_t` domain.

(BZ#[1365944](#))

The SCAP password length requirement is ignored in the kickstart installation

The interactive kickstart installation does not enforce the password length check defined by the SCAP rule and accepts shorter root passwords. To work around this problem, use the `--strict` option with the `pwpolicy root` command in the kickstart file. (BZ#1372791)

rhnsd.pid is writable by group and others

In Red Hat Enterprise Linux 7.4, the default permissions of the `/var/run/rhnsd.pid` file are set to `-rw-rw-rw-..`. This setting is not secure. To work around this problem, change the permissions of this file to be writable only by the owner:

```
# chmod go-w /var/run/rhnsd.pid
```

(BZ#1480306)

CHAPTER 63. STORAGE

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

While RAID logical volumes and thinly provisioned logical volumes can be used in a cluster when activated exclusively, there is currently no support for thin provisioning on top of RAID in a cluster. This is the case even if the combination is activated exclusively. Currently this combination is only supported in LVM's single machine non-clustered mode. (BZ#[1014758](#))

Anaconda installation can fail when LVM or md device has metadata from a previous install

During Red Hat Enterprise Linux 7 installation on a machine where a disk to be multipathed already starts with LVM or md metadata on it from a previous install, multipath will not get set up on the device, and LVM/md will get set up on one of the path devices while Anaconda is starting up. This can create problems with Anaconda, and cause the installation to fail. The workaround for this issue is to add `mpath.wwid=<WWID>` to the kernel command line when booting up for the installation. `<WWID>` is the wwid of the device that multipath should claim. This value is also the same as the `ID_SERIAL` udev database value for scsi devices and `ID_UID` for DASD devices. (BZ#1378714)

CHAPTER 64. SYSTEM AND SUBSCRIPTION MANAGEMENT

System upgrade may cause Yum to install unneeded 32-bit packages if rdma-core is installed

In Red Hat Enterprise Linux 7.4, the `rdma-core.noarch` packages are obsoleted by `rdma-core.i686` and `rdma-core.x86_64`. During a system upgrade, Yum replaces the original package with both of the new packages, and installs any required dependencies. This means that the 32-bit package, as well a potentially large amount of its 32-bit dependencies, is installed by default, even if not required.

To work around this problem, you can either use the `yum update` command with the `--exclude='*.i686'` option, or you can use `yum remove rdma-core.i686` after the upgrade to remove the 32-bit package. (BZ#1458338)

CHAPTER 65. VIRTUALIZATION

Booting OVMF guests fails

Attempting to boot a guest virtual machine that uses the Open Virtual Machine Firmware (OVMF) on a Red Hat Enterprise Linux host using the `qemu-kvm` package currently fails, with the guest becoming unresponsive and displaying a blank screen. (BZ#1174132)

Bridge creation with `virsh iface-bridge` fails

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

For more information, see the Red Hat Knowledgebase: <https://access.redhat.com/solutions/2792701> (BZ#1100588)

Guests sometimes fail to boot on ESXi 5.5

When running Red Hat Enterprise Linux 7 guests with 12 GB RAM or above on a VMware ESXi 5.5 hypervisor, certain components currently initialize with incorrect memory type range register (MTRR) values or incorrectly reconfigure MTRR values across boots. This sometimes causes the guest kernel to panic or the guest to become unresponsive during boot.

To work around this problem, add the `disable_mtrr_trim` option to the guest's kernel command line, which enables the guest to continue booting when MTRRs are configured incorrectly. Note that with this option, the guest prints **WARNING: BIOS bug** messages during boot, which you can safely ignore. (BZ#1429792)

The STIG for Red Hat Virtualization Hypervisor profile is not displayed in Anaconda

The `oscap-anaconda-addon` module is currently not able to properly parse the **STIG for Red Hat Virtualization Hypervisor** security hardening profile. As a consequence, the profile's name is shown as **DISA STIG for Red Hat Enterprise Linux 7 or United States Government Configuration Baseline (USGCB / STIG) - DRAFT** in the Anaconda interface selection. However, this is only a display problem, and you can safely use the **DISA STIG for Red Hat Enterprise Linux 7** profile instead of the **STIG for Red Hat Virtualization Hypervisor** profile. (BZ#1437106)

APPENDIX A. COMPONENT VERSIONS

This appendix provides a list of key components and their versions in the Red Hat Enterprise Linux 7.4 release.

Table A.1. Component Versions

Component	Version
Kernel	3.10.0-693
QLogic qla2xxx driver	8.07.00.38.07.4-k1
QLogic qla4xxx driver	5.04.00.00.07.02-k0
Emulex lpfc driver	0:11.2.0.6
iSCSI initiator utils	iscsi-initiator-utils-6.2.0.874-4
DM-Multipath	device-mapper-multipath-0.4.9-111
LVM	lvm2-2.02.171-8

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# 1394000 , BZ# 1395940 , BZ# 1425907	BZ# 1378209		
Doc-administration-guide	BZ# 1426286 , BZ# 1426289			
Doc-release-notes	BZ# 1426275 , BZ# 1426278 , BZ# 1426283 , BZ# 1436973			
NetworkManager	BZ# 1337997 , BZ# 1353612 , BZ# 1373698 , BZ# 1394344 , BZ# 1394579 , BZ# 1398934 , BZ# 1404594 , BZ# 1404598 , BZ# 1414103 , BZ# 1420708	BZ# 1391170 , BZ# 1393997		
aide		BZ# 1377215		
anaconda	BZ# 663099 , BZ# 1113207 , BZ# 1131247 , BZ# 1255659 , BZ# 1315160 , BZ# 1332316 , BZ# 1366935 , BZ# 1377233 , BZ# 1391724 , BZ# 1412538	BZ# 1317370 , BZ# 1327439 , BZ# 1356975 , BZ# 1358778 , BZ# 1373360 , BZ# 1380224 , BZ# 1380277 , BZ# 1404158 , BZ# 1412022 , BZ# 1441337		BZ# 1378714
ansible			BZ# 1313263	
audit	BZ# 1381601			
authconfig	BZ# 1334449 , BZ# 1378943			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
autofs	BZ# 1367576 , BZ# 1382093	BZ# 1101782 , BZ# 1383194 , BZ# 1383910 , BZ# 1420574 , BZ# 1420584 , BZ# 1320588		
bind	BZ# 1388534 , BZ# 1393886			
bind-dyndb-ldap	BZ# 1393889			
binutils	BZ# 1366052	BZ# 1326710 , BZ# 1406498		
bison	BZ# 1306000			
booth	BZ# 1302087			
ca-certificates	BZ# 1444414			
chrony	BZ# 1387223			
chrtpath		BZ# 1271380		
clevis	BZ# 1300697			
cloud-init	BZ# 1427280			
clutter	BZ# 1387424			
crash	BZ# 1368711 , BZ# 1384944 , BZ# 1393534			
criu			BZ# 1400230	
custodia			BZ# 1403214	
cyrus-sasl		BZ# 1421663		
dbxtool	BZ# 1078990			
dconf-editor	BZ# 1388931			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
device-mapper-multipath	BZ# 1169168 , BZ# 1279355 , BZ# 1359510 , BZ# 1362409 , BZ# 1368211 , BZ# 1372032 , BZ# 1394059 , BZ# 1406226 , BZ# 1416569 , BZ# 1430097	BZ# 1239173 , BZ# 1362120 , BZ# 1380602 , BZ# 1402092 , BZ# 1403552 , BZ# 1431562		
dhcp	BZ# 1374119			
distribution				BZ# 1062656
dmidecode	BZ# 1385884			
dnsmasq	BZ# 1375527 , BZ# 1375569			
ecj	BZ# 1379855			
elfutils	BZ# 1400302			
empathy		BZ# 1386616		
ethtool	BZ# 1402701			
fcoe-utils		BZ# 1384707		
firefox				BZ# 1455798
firewalld	BZ# 1006225 , BZ# 1409544 , BZ# 1419058	BZ# 1401978		
flatpak				BZ# 1476905
genwqe-tools	BZ# 1275663			
gfs2-utils	BZ# 1413684			
ghostscript		BZ# 1390847 , BZ# 1411725 , BZ# 1424752		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
git		BZ# 1369173		
glibc	BZ# 841653 , BZ# 1298975 , BZ# 1320947 , BZ# 1421155	BZ# 1324568 , BZ# 1326739		
glusterfs				BZ# 1409773
gnome-initial-setup		BZ# 1226819		
gnome-packagekit				BZ# 1387181
gnome-shell	BZ# 1383353			
gnome-software				BZ# 1434477 , BZ# 1464139
gnu-efi	BZ# 1310782			
gnutls	BZ# 1399232			
grep	BZ# 1297441			
grub2				BZ# 1154226
gstreamer1-plugins-good				BZ# 1451211
http-parser	BZ# 1393819			
hwdata	BZ# 1386133			
initial-setup		BZ# 1378082		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
initscripts	BZ# 1260552 , BZ# 1428935	BZ# 1278521 , BZ# 1367554 , BZ# 1369790 , BZ# 1374837 , BZ# 1385272 , BZ# 1392766 , BZ# 1394191 , BZ# 1398671 , BZ# 1398678 , BZ# 1398679 , BZ# 1398683 , BZ# 1398686 , BZ# 1406254 , BZ# 1408219 , BZ# 1428574 , BZ# 1434075		
intel-cmt-cat	BZ# 1315489			
ipa	BZ# 872671 , BZ# 1125174 , BZ# 1200767 , BZ# 1366572 , BZ# 1402959 , BZ# 1404750 , BZ# 1409628 , BZ# 1459153		BZ# 1115294 , BZ# 1298286	BZ# 1455946
ipa-server-docker			BZ# 1405325	
iperf3	BZ# 913329			
ipmitool	BZ# 1398658			
iproute	BZ# 1063934 , BZ# 1422629	BZ# 1375215		
iprutils	BZ# 1384382			
jansson	BZ# 1389805			
java-1.7.0-openjdk	BZ# 1373986			
java-1.8.0-openjdk			BZ# 1400306	
kernel	BZ# 437984 , BZ# 950243 , BZ# 1072503 ,	BZ# 1084802 , BZ# 1213119 , BZ# 1217546 ,	BZ# 916382 , BZ# 947163 , BZ# 1109348 ,	BZ# 1377857 , BZ# 1380447 , BZ# 1429792 ,

Component	BZ# 1138782 , BZ# 1155732 , BZ# 1241990 , BZ# 1270982 , BZ# 1273401 , BZ# 1288964 , BZ# 1297841 , BZ# 1297929 , BZ# 1298643 , BZ# 1299527 , BZ# 1302147 , BZ# 1306396 , BZ# 1306453 , BZ# 1308632 , BZ# 1314179 , BZ# 1326309 , BZ# 1326318 , BZ# 1326353 , BZ# 1330457 , BZ# 1339127 , BZ# 1339791 , BZ# 1340238 , BZ# 1346348 , BZ# 1352289 , BZ# 1355919 , BZ# 1356122 , BZ# 1357491 , BZ# 1365002 , BZ# 1366564 , BZ# 1369158 , BZ# 1373606 , BZ# 1373971 , BZ# 1374498 , BZ# 1377710 , BZ# 1377767 , BZ# 1379590 , BZ# 1382101 , BZ# 1382494 , BZ# 1382500 , BZ# 1382504 , BZ# 1382508 , BZ# 1382849 , BZ# 1383280 , BZ# 1383827 , BZ# 1383834 , BZ# 1384456 , BZ# 1384648 , BZ# 1385026 , BZ# 1385757 , BZ# 1388467 , BZ# 1388646 , BZ# 1388716 , BZ# 1391219 , BZ# 1391243 , BZ# 1391413 , BZ# 1391668 ,	BZ# 1324918 , BZ# 1331098 , BZ# 1353218 , BZ# 1363661 , BZ# 1370638 , BZ# 1379787 , BZ# 1385149 , BZ# 1386923 , BZ# 1387485 , BZ# 1406885 , BZ# 1408330 , BZ# 1412898	BZ# 1111712 , BZ# 1205497 , BZ# 1208277 , BZ# 1230959 , BZ# 1274456 , BZ# 1274459 , BZ# 1299662 , BZ# 1305092 , BZ# 1348508 , BZ# 1349668 , BZ# 1350553 , BZ# 1393375 , BZ# 1414957 , BZ# 1449762 , BZ# 1460849	BZ# 1431964 , BZ# 1431242 , BZ# 1458354
-----------	--	--	---	---

Component	New Features BZ#1394197, BZ#1400501, BZ#1401797, BZ#1402102, BZ#1406197, BZ#1416924, BZ#1432218, BZ#1432897, BZ#1383489	Notable Bug Fixes	Technology Previews	Known Issues
kernel-aarch64				BZ#1218374
kernel-rt	BZ#1391779	BZ#1443711	BZ#1297061	
kexec-tools	BZ#1384945			
keycloak-httpd-client-install	BZ#1401781			
libcgroup		BZ#1406927		
libdb		BZ#1277887		
libfastjson	BZ#1395145			
libguestfs	BZ#1233093, BZ#1359086, BZ#1362649, BZ#1367738, BZ#1400205, BZ#1404182	BZ#1161019, BZ#1265588, BZ#1311890, BZ#1354507, BZ#1374232, BZ#1374405, BZ#1383517, BZ#1392798, BZ#1401474, BZ#1402301, BZ#1431579	BZ#1387213, BZ#1441197	
libica	BZ#1391558			
libnfsidmap	BZ#980925			
libnftnl		BZ#1418967	BZ#1332585	
libreoffice				BZ#1466164
libreswan	BZ#1324458, BZ#1399883			
librtas	BZ#1380656			
libseccomp	BZ#1425007			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
libstoragemgmt	BZ# 1403142		BZ# 1119909	
libusnic_verbs			BZ# 916384	
libvirt	BZ# 1349696 , BZ# 1382640 , BZ# 1414627		BZ# 1283251	
libwacom	BZ# 1342990			BZ# 1458351
linuxptp	BZ# 1359311			
logrotate	BZ# 1381719			
lorax	BZ# 1310775 , BZ# 1430483			BZ# 1341280
lshw	BZ# 1368704			
lvm2	BZ# 1189108 , BZ# 1191935 , BZ# 1346280 , BZ# 1366296 , BZ# 1378956 , BZ# 1394048 , BZ# 1436748 , BZ# 1442992	BZ# 1380521 , BZ# 1380532 , BZ# 1382688 , BZ# 1386184 , BZ# 1434054		BZ# 1014758
mariadb		BZ# 1356897		
mdadm	BZ# 1380017			
memkind	BZ# 1384549			
mod_nss	BZ# 1382102 , BZ# 1392582			
mutt	BZ# 1388511	BZ# 1388512		
mutter				BZ# 1393951
nautilus				BZ# 1474852
net-snmp		BZ# 1286693 , BZ# 1324306		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
netcf				BZ#1100588
nfs-utils	BZ#1375259, BZ#1418041			BZ#1450447
nss	BZ#1309781, BZ#1316546, BZ#1444413	BZ#1220573		BZ#1401809
nss-softokn	BZ#1369055			
nvme-cli	BZ#1382119			
nvmeccli	BZ#1383837			
opencryptoki	BZ#1391559			
openldap	BZ#1386365, BZ#1428740			
opensc	BZ#1081088			
openscap	BZ#1363826	BZ#1420038, BZ#1440192, BZ#1447341		
openssh	BZ#1322911, BZ#1341754	BZ#1418062		
openssl	BZ#1276310			
openssl-ibmca	BZ#1274385			
openvswitch	BZ#1368043, BZ#1390938			
openwsman	BZ#1190689			
oprofile		BZ#1380809		
oscap-anaconda-addon				BZ#1372791, BZ#1437106, BZ#1462647

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
other	BZ# 1432080 , BZ# 1444937 , BZ# 1457907 , BZ# 1459948 , BZ# 1467260	BZ# 1408694	BZ# 1062759 , BZ# 1072107 , BZ# 1259547 , BZ# 1464377 , BZ# 1467338 , BZ# 1477977	BZ# 1174132 , BZ# 1458338 , BZ# 1463600
OVMF			BZ# 653382	
pacemaker	BZ# 1289662	BZ# 1388489		
pcp	BZ# 1422263 , BZ# 1423020			
pcrc		BZ# 1400267		BZ# 1290432
pcs	BZ# 1158805 , BZ# 1165821 , BZ# 1176018 , BZ# 1261116 , BZ# 1303969 , BZ# 1362493 , BZ# 1373614 , BZ# 1413958	BZ# 1386114 , BZ# 1378107	BZ# 1433016	
perl-IO-Socket-SSL	BZ# 1335035			
perl-Net-SSLeay	BZ# 1335028			
perl-Perl4-CoreLibs	BZ# 1366724			
perl-local-lib		BZ# 1122993		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
pki-core	BZ# 1303683 , BZ# 1305993 , BZ# 1325071 , BZ# 1388622 , BZ# 1391737 , BZ# 1392068 , BZ# 1409946 , BZ# 1413132 , BZ# 1426754 , BZ# 1445535 , BZ# 1447144 , BZ# 1450143 , BZ# 1458055	BZ# 1238684 , BZ# 1246635 , BZ# 1249400 , BZ# 1282504 , BZ# 1330800 , BZ# 1372052 , BZ# 1376226 , BZ# 1376488 , BZ# 1378275 , BZ# 1378277 , BZ# 1381084 , BZ# 1382066 , BZ# 1385208 , BZ# 1386303 , BZ# 1395817 , BZ# 1397200 , BZ# 1400149 , BZ# 1404881 , BZ# 1411428 , BZ# 1412681 , BZ# 1413136 , BZ# 1445088 , BZ# 1446364 , BZ# 1447762 , BZ# 1452250 , BZ# 1452344 , BZ# 1454450 , BZ# 1454471 , BZ# 1458429		BZ# 1256901
procps-ng		BZ# 1373246		
psacct	BZ# 1255183			
pykickstart				BZ# 1452770
python	BZ# 1219110			
python-blivet		BZ# 1214407 , BZ# 1327463		
python-tornado	BZ# 1158617			
qemu-kvm			BZ# 1103193	
rdma-core	BZ# 1404035			
rear	BZ# 1355667	BZ# 1343119		

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
resource-agents	BZ# 1077888 , BZ# 1336847 , BZ# 1430304			
rhino	BZ# 1350331			
rhnsd				BZ# 1480306
rpm		BZ# 1378307	BZ# 1278924	
rsyslog	BZ# 1313490 , BZ# 1431616			
ruby	BZ# 1397390	BZ# 1308992		
rubygem-abrt	BZ# 1418750			
samba	BZ# 1391954			
sapconf		BZ# 1391881		
sbd	BZ# 1413951			
sblim-cmpi-fsvol		BZ# 1136116		
scap-security-guide	BZ# 1404392 , BZ# 1410914	BZ# 1450731		BZ# 1448952 , BZ# 1464899 , BZ# 1465677
seabios		BZ# 1020622		
selinux-policy		BZ# 1368057 , BZ# 1386916		BZ# 1365944
sendmail	BZ# 1124827			
sg3_utils		BZ# 1380744		
shim	BZ# 1310766			
shim-signed	BZ# 1310764			
sos	BZ# 1414879			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
sssd	BZ# 1214491 , BZ# 1311056 , BZ# 1330196 , BZ# 1340711 , BZ# 1396012 , BZ# 1414023 , BZ# 1425891		BZ# 1068725	BZ# 1446101 , BZ# 1447945 , BZ# 1448094 , BZ# 1460689 , BZ# 1462769
strace		BZ# 1377847		
strongimcv			BZ# 755087	
sudo				BZ# 1293306
system-config-language		BZ# 1304223		
systemd		BZ# 1353028	BZ# 1284974	
systemtap	BZ# 1398393			
tar	BZ# 1350640	BZ# 1184697 , BZ# 1319820 , BZ# 1341786		
targetcli	BZ# 1243410			
targetd	BZ# 1162381			
tboot	BZ# 1384210			
tcpdump	BZ# 1292056 , BZ# 1422473			
tcsh		BZ# 1388426		
telnet	BZ# 1367415			
tpm2-tss	BZ# 1275027			
tss2	BZ# 1384452			
tuned	BZ# 1388454 , BZ# 1414098			
unbound	BZ# 1382383			

Component	New Features	Notable Bug Fixes	Technology Previews	Known Issues
usbguard	BZ# 1395615		BZ# 1467369	
valgrind	BZ# 1391217			
virt-who	BZ# 1299643 , BZ# 1369107 , BZ# 1405967 , BZ# 1426058 , BZ# 1436811			
wget	BZ# 1439811			
wpa_supplicant	BZ# 1404793			
xorg-x11-drv-libinput	BZ# 1413811			
xorg-x11-drv-qxl				BZ# 1428340
xorg-x11-server	BZ# 1404868			BZ# 1185690
yelp				BZ# 1443179
yp-tools		BZ# 1401432		
ypbind		BZ# 1217435 , BZ# 1382804		
yum	BZ# 1343690	BZ# 1352585 , BZ# 1370134		
yum-utils		BZ# 1406891		

APPENDIX C. REVISION HISTORY

Revision 0.4-1	Thu Mar 22 2018	Lenka Špačková
Fixed the openldap-servers package name (Deprecated Functionality).		
Revision 0.4-0	Fri Mar 16 2018	Lenka Špačková
Added a pcs-related bug fix (Clustering).		
Revision 0.3-9	Mon Feb 19 2018	Mirek Jahoda
The incorrectly placed TPM-related features moved to the Technnology Preview section.		
Revision 0.3-8	Tue Feb 06 2018	Lenka Špačková
Added a missing Technology Preview - OVMF (Virtualization).		
Added information regarding deprecation of containers using the <code>libvirt-lxc</code> tooling.		
Revision 0.3-7	Wed Jan 17 2018	Lenka Špačková
Updated the FCoE deprecation notice.		
Revision 0.3-6	Wed Jan 10 2018	Lenka Špačková
Changed the status of Device DAX for NVDIMM devices from Technology Preview to fully supported (Storage).		
Revision 0.3-5	Thu Dec 14 2017	Lenka Špačková
Unified the structure of deprecated drivers.		
Revision 0.3-4	Tue Dec 12 2017	Lenka Špačková
Updated deprecated adapters from the <code>qla2xxx</code> driver.		
Revision 0.3-3	Wed Nov 22 2017	Lenka Špačková
Added information regarding <code>pam_krb5</code> to <code>sssd</code> migration (Deprecated Functionality).		
Revision 0.3-2	Wed Nov 15 2017	Lenka Špačková
Fixed a typo.		
Revision 0.3-1	Tue Oct 31 2017	Lenka Špačková
Added an LVM-related bug fix description (Storage).		
Revision 0.3-3	Mon Oct 30 2017	Lenka Špačková
Added an <code>autofs</code> bug fix description (File Systems).		
Added information on changes in the <code>ld</code> linker behavior to Deprecated Functionality.		
Revision 0.3-2	Wed Sep 13 2017	Lenka Špačková
Added information regarding limited support for visuals in the Xorg server.		
Revision 0.3-1	Mon Sep 11 2017	Lenka Špačková
Added CUIR enhanced scope detection to Technology Previews (Kernel).		
Updated openssh rebase description in New Features (Security).		
Revision 0.3-0	Mon Sep 04 2017	Lenka Špačková
Added two known issues (Security, Desktop).		
Revision 0.2-9	Mon Aug 21 2017	Lenka Špačková
Added <code>tcp_wrappers</code> to Deprecated Functionality.		
Revision 0.2-8	Tue Aug 15 2017	Lenka Špačková
Added several new features and a known issue.		
Revision 0.2-7	Mon Aug 14 2017	Lenka Špačková

Removed a duplicate note.

Revision 0.2-6 Updated several known issues.	Thu Aug 10 2017	Lenka Špačková
Revision 0.2-5 Added two known issues.	Tue Aug 08 2017	Lenka Špačková
Revision 0.2-4 Updated FCoE deprecation notice. Minor updates and additions.	Mon Aug 07 2017	Lenka Špačková
Revision 0.2-3 Moved several new features from Virtualization to System and Subscription Management.	Fri Aug 04 2017	Lenka Špačková
Revision 0.2-2 Updated information on Btrfs ; it is now both in the Technology Previews and Deprecated Functionality parts. Minor updates and additions.	Thu Aug 03 2017	Lenka Špačková
Revision 0.2-1 Release of the Red Hat Enterprise Linux 7.4 Release Notes.	Tue Aug 01 2017	Lenka Špačková
Revision 0.0-4 Release of the Red Hat Enterprise Linux 7.4 Beta Release Notes.	Tue May 23 2017	Lenka Špačková