



Red Hat Enterprise Linux 6

6.5 Technical Notes

Detailed notes on the changes implemented in Red Hat Enterprise Linux 6.5
Edition 5

Red Hat Enterprise Linux 6 6.5 Technical Notes

Detailed notes on the changes implemented in Red Hat Enterprise Linux 6.5
Edition 5

Red Hat Customer Content Services

Legal Notice

Copyright © 2013-2016 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). 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 Red Hat Enterprise Linux 6.5 Technical Notes list and document the changes made to the Red Hat Enterprise Linux 6 operating system and its accompanying applications between Red Hat Enterprise Linux 6.4 and minor release Red Hat Enterprise Linux 6.5.

Table of Contents

PREFACE	7
CHAPTER 1. RED HAT ENTERPRISE LINUX 6.5 INTERNATIONAL LANGUAGES	8
CHAPTER 2. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS	10
CHAPTER 3. DEVICE DRIVERS	13
Storage Drivers	13
Network Drivers	13
Miscellaneous Drivers	13
CHAPTER 4. TECHNOLOGY PREVIEWS	15
4.1. STORAGE AND FILE SYSTEMS	15
4.2. NETWORKING	16
4.3. CLUSTERING AND HIGH AVAILABILITY	17
4.4. AUTHENTICATION	18
4.5. SECURITY	18
4.6. DEVICES	18
4.7. KERNEL	18
CHAPTER 5. DEPRECATED FUNCTIONALITY	20
CHAPTER 6. KNOWN ISSUES	22
6.1. INSTALLATION	22
6.2. ENTITLEMENT	23
6.3. DEPLOYMENT	24
6.4. VIRTUALIZATION	24
6.5. STORAGE AND FILE SYSTEMS	28
6.6. NETWORKING	31
6.7. SECURITY	35
6.8. CLUSTERING	35
6.9. AUTHENTICATION	36
6.10. DEVICES	43
6.11. KERNEL	44
6.12. DESKTOP	49
6.13. TOOLS	50
6.14. DOCUMENTATION	51
CHAPTER 7. NEW PACKAGES	53
7.1. RHEA-2013:1625 — NEW PACKAGES: FREERDP	53
7.2. RHBA-2013:1607 — NEW PACKAGES: GCC-LIBRARIES	53
7.3. RHEA-2013:1728 — NEW PACKAGES: OPENHPI32	53
7.4. RHEA-2013:1626 — NEW PACKAGES: P11-KIT	53
7.5. RHEA-2013:1621 — NEW PACKAGE: PS_MEM	54
7.6. RHEA-2013:1642 — NEW PACKAGES: REDHAT-SUPPORT-LIB-PYTHON AND REDHAT-SUPPORT-TOOL	54
7.7. RHEA-2013:1686 — NEW PACKAGE: SAPCONF	55
7.8. RHEA-2013:1731 — NEW PACKAGES: SNAPPY	55
7.9. RHEA-2013:1622 — NEW PACKAGES: XORG-X11-GLAMOR	55
CHAPTER 8. UPDATED PACKAGES	56
8.1. ABRT	56
8.2. ANACONDA	57
8.3. ARPTABLES_JF	62

8.4. AUGEAS	63
8.5. AUTOFS	64
8.6. BATIK	66
8.7. BFA-FIRMWARE	66
8.8. BIND-DYNDB-LDAP	67
8.9. BIOSDEVNAME	68
8.10. BOOST	68
8.11. BUSYBOX	69
8.12. CA-CERTIFICATES	70
8.13. CIFS-UTILS	70
8.14. CJKUNI-FONTS	70
8.15. CLUSTER AND GFS2-UTILS	71
8.16. CLUSTERMON	73
8.17. COMPAT-OPENMPI	74
8.18. CONMAN	74
8.19. COOLKEY	75
8.20. COREUTILS	76
8.21. COROSYNC	78
8.22. CPUPOWERUTILS	80
8.23. CRASH	82
8.24. CRASH-GCORE-COMMAND	82
8.25. CREATEREPO	83
8.26. CRONIE	83
8.27. CVS	85
8.28. DEVICE-MAPPER-MULTIPATH	86
8.29. DEVICE-MAPPER-PERSISTENT-DATA	89
8.30. DHCP	89
8.31. DOVECOT	90
8.32. DRACUT	91
8.33. E2FSPROGS	94
8.34. EFIBOOTMGR	95
8.35. EMACS	95
8.36. ENVIRONMENT-MODULES	96
8.37. ESC	97
8.38. EVOLUTION	97
8.39. FCOE-TARGET-UTILS	98
8.40. FCOE-UTILS	99
8.41. FEBOOTSTRAP	101
8.42. FENCE-AGENTS	101
8.43. FENCE-VIRT	103
8.44. FIRSTBOOT	103
8.45. FOOMATIC	104
8.46. FPRINTD	104
8.47. FREEIPMI	105
8.48. FTP	105
8.49. GCC	106
8.50. GDM	107
8.51. GEGL	108
8.52. GHOSTSCRIPT	109
8.53. GLIB2	109
8.54. GLIBC	110
8.55. GLUSTERFS	114
8.56. GNOME-SCREENSAVER	115

8.57. GPXE	115
8.58. GREP	116
8.59. GRUB	116
8.60. GRUBBY	118
8.61. GTK2	118
8.62. HAPROXY	119
8.63. HDPARM	120
8.64. HSQLDB	121
8.65. HWDATA	121
8.66. HYPERVKVPD	122
8.67. IBUS-HANGUL	123
8.68. ICEDTEA-WEB	123
8.69. INITSCRIPTS	124
8.70. IOTOP	126
8.71. IPA	126
8.72. IPMITOOL	130
8.73. IPROUTE	131
8.74. IPTABLES	132
8.75. IPVSADM	133
8.76. IRQBALANCE	134
8.77. ISCSI-INITIATOR-UTILS	134
8.78. IW	136
8.79. JAVA-1.6.0-OPENJDK	136
8.80. JAVA-1.7.0-OPENJDK	136
8.81. KDE-SETTINGS	137
8.82. KERNEL	138
8.83. KEXEC-TOOLS	191
8.84. KSH	194
8.85. LEDMON	196
8.86. LIBXCURSOR	196
8.87. LIBCGROUP	197
8.88. LIBDRM	198
8.89. LIBGUESTFS	198
8.90. LIBIBVERBS-ROCEE	201
8.91. LIBKSBA	201
8.92. LIBNL	202
8.93. LIBPCAP	202
8.94. LIBQB	203
8.95. LIBREOFFICE	204
8.96. LIBRTAS	205
8.97. LIBTEVENT	206
8.98. LIBVIRT	207
8.99. LIBVIRT-CIM	218
8.100. LIBVIRT-SNMP	219
8.101. LIBWACOM	219
8.102. LIBXML2	220
8.103. LINUXPTP	220
8.104. LKSCTP-TOOLS	221
8.105. LOGROTATE	222
8.106. LOGWATCH	223
8.107. LUCI	224
8.108. LVM2	227
8.109. MAILX	234

8.110. MAN-PAGES-FR	234
8.111. MAN-PAGES-JA	235
8.112. MAN-PAGES-OVERRIDES	235
8.113. MCELOG	238
8.114. MDADM	238
8.115. MESA	240
8.116. MICROCODE_CTL	240
8.117. MOBILE-BROADBAND-PROVIDER-INFO	241
8.118. MOD_AUTH_KERB	241
8.119. MODEMMANAGER	242
8.120. MYSQL	242
8.121. NET-SNMP	243
8.122. NETCF	246
8.123. NETWORKMANAGER	246
8.124. NFS-UTILS	249
8.125. NMAP	250
8.126. NSS AND NSPR	251
8.127. NTP	253
8.128. NUMACTL	254
8.129. NUMAD	255
8.130. OPENCRIPTOKI	256
8.131. OPENCV	256
8.132. OPENHPI	257
8.133. OPENSCAP	258
8.134. OPENSSSH	258
8.135. OPENSSSL	260
8.136. OPENSWAN	262
8.137. PACEMAKER	265
8.138. PAM	267
8.139. PAPI	268
8.140. PARTED	269
8.141. PCS	270
8.142. PERL	272
8.143. PERL-CGI-SESSION	273
8.144. PERL-CONFIG-GENERAL	273
8.145. PERL-DATETIME	275
8.146. PERL-MAKEFILE-PARSER	276
8.147. PERL-NET-DNS	276
8.148. PERL-SOCKET6	277
8.149. PERL-TEST-MEMORY-CYCLE	277
8.150. PERL-TEST-MOCKOBJECT	277
8.151. PERL-XML-DUMPER	278
8.152. PHP	278
8.153. PIRANHA	280
8.154. 389-DS-BASE	280
8.155. PKI-CORE	287
8.156. POLICYCOREUTILS	287
8.157. POWERTOP	289
8.158. PYKICKSTART	289
8.159. PYPARTED	290
8.160. PYTHON	291
8.161. PYTHON-BEAKER	292
8.162. PYTHON-ETHOTOOL	293

8.163. PYTHON-URLGRABBER	293
8.164. PYTHON-URWID	294
8.165. PYTHON-VIRTINST	294
8.166. PYTHON-WEBERROR	296
8.167. QEMU-KVM	296
8.168. QL2400-FIRMWARE	299
8.169. QL2500-FIRMWARE	299
8.170. QUOTA	299
8.171. RDESKTOP	300
8.172. RDMA STACK	301
8.173. READAHEAD	302
8.174. REDHAT-INDEXHTML	303
8.175. REDHAT-RELEASE	303
8.176. RED HAT ENTERPRISE LINUX 6.5 RELEASE NOTES	303
8.177. RESOURCE-AGENTS	304
8.178. RGMANAGER	307
8.179. RHEL-GUEST-IMAGE	308
8.180. RHN-CLIENT-TOOLS	309
8.181. RHNLIB	310
8.182. RICCI	310
8.183. RP-PPPOE	311
8.184. RPM	312
8.185. RPMLINT	313
8.186. RSYSLOG	313
8.187. RUBYGEMS	314
8.188. S390UTILS	315
8.189. SAMBA	317
8.190. SAMBA4	320
8.191. SANLOCK	320
8.192. SBLIM-CMPI-FSVOL	321
8.193. SBLIM-SFCC	321
8.194. SBLIM-WBEMCLI	322
8.195. SCL-UTILS	322
8.196. SCSI-TARGET-UTILS	323
8.197. SEABIOS	324
8.198. SELINUX-POLICY	325
8.199. SETUPTOOL	334
8.200. SG3_UTILS	334
8.201. SLAPI-NIS	335
8.202. SOS	335
8.203. SPICE-GTK	337
8.204. SPICE-PROTOCOL	339
8.205. SPICE-SERVER	339
8.206. SPICE-VDAGENT	342
8.207. SPICE-XPI	343
8.208. SSSD	344
8.209. SUBSCRIPTION-MANAGER	348
8.210. SUDO	351
8.211. SUITESPARSE	353
8.212. SYSSTAT	353
8.213. SYSTEM-CONFIG-DATE	354
8.214. SYSTEM-CONFIG-KEYBOARD	355
8.215. SYSTEM-CONFIG-LVM	355

8.216. SYSTEM-CONFIG-USERS-DOCS	355
8.217. SYSTEMTAP	356
8.218. SYSVINIT	358
8.219. TALK	358
8.220. TBOOT	358
8.221. TOMCAT6	359
8.222. TUNED	360
8.223. UDEV	361
8.224. UTIL-LINUX-NG	363
8.225. VHOSTMD	364
8.226. VIRT-MANAGER	365
8.227. VIRT-P2V	366
8.228. VIRT-V2V	366
8.229. VIRT-VIEWER	368
8.230. VIRT-WHO	370
8.231. VIRTIO-WIN	371
8.232. WATCHDOG	375
8.233. WEBKITGTK	376
8.234. WIRESHARK	376
8.235. XFSPROGS	378
8.236. XMLRPC-C	379
8.237. XORG-X11-DRV-ATI	379
8.238. XORG-X11-DRV-INTEL	380
8.239. XORG-X11-DRV-MGA	380
8.240. XORG-X11-DRV-NOUVEAU	381
8.241. XORG-X11-DRV-QXL	381
8.242. XORG-X11-DRV-SYNAPTICS	382
8.243. XORG-X11-DRV-WACOM	382
8.244. XORG-X11-SERVER	383
8.245. XORG-X11-XINIT	383
8.246. YABOOT	384
8.247. YUM-RHN-PLUGIN	385
8.248. ZSH	385
APPENDIX A. REVISION HISTORY	387

PREFACE

The *Red Hat Enterprise Linux 6.5 Technical Notes* list and document the changes made to the Red Hat Enterprise Linux 6 operating system and its accompanying applications between minor release Red Hat Enterprise Linux 6.4 and minor release Red Hat Enterprise Linux 6.5.

For system administrators and others planning Red Hat Enterprise Linux 6.5 upgrades and deployments, the Technical Notes provide a single, organized record of the bugs fixed in, features added to, and Technology Previews included with this new release of Red Hat Enterprise Linux.

For auditors and compliance officers, the *Red Hat Enterprise Linux 6.5 Technical Notes* provide a single, organized source for change tracking and compliance testing.

For every user, the *Red Hat Enterprise Linux 6.5 Technical Notes* provide details of what has changed in this new release.



NOTE

The [Package Manifest](#) is available as a separate document.

CHAPTER 1. RED HAT ENTERPRISE LINUX 6.5 INTERNATIONAL LANGUAGES

Red Hat Enterprise Linux 6.5 supports installation of multiple languages and changing of languages based on your requirements.

The following languages are supported in Red Hat Enterprise Linux 6.5:

- East Asian Languages - Japanese, Korean, Simplified Chinese, and Traditional Chinese
- European Languages - English, German, Spanish, French, Portuguese Brazilian, and Russian,
- Indic Languages - Assamese, Bengali, Gujarati, Hindi, Kannada, Malayalam, Marathi, Oriya, Punjabi, Tamil, and Telugu

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

Table 1.1. Red Hat Enterprise Linux 6 International Languages

Territory	Language	Locale	Fonts	Package Names
China	Simplified Chinese	zh_CN.UTF-8	AR PL (ShanHeiSun and Zenkai) Uni	fonts-chinese, scim-pinyin, scim-tables
Japan	Japanese	ja_JP.UTF-8	Sazanami (Gothic and Mincho)	fonts-japanese, scim-anthy
Korea	Hangul	ko_KR.UTF-8	Baekmuk (Batang, Dotum, Gulim, Headline)	fonts-korean, scim-hangul
Taiwan	Traditional Chinese	zh_TW.UTF-8	AR PL (ShanHeiSun and Zenkai) Uni	fonts-chinese, scim-chewing, scim-tables
Brazil	Portuguese	pt_BR.UTF-8	standard latin fonts	
France	French	fr_FR.UTF-8	standard latin fonts	
Germany	German	de_DE.UTF-8	standard latin fonts	
Italy	Italy	it_IT.UTF-8	standard latin fonts	
Russia	Russian	ru_RU.UTF-8	KOI8-R, fonts-KOI8-R-100dpi, fonts-KOI8-R-75dpi and xorg-x11-fonts-cyrillic	fonts-KOI8-R, fonts-KOI8-R-100dpi, fonts-KOI8-R-75dpi, xorg-x11-fonts-cyrillic

Territory	Language	Locale	Fonts	Package Names
Spain	Spanish	es_ES.UTF-8	standard latin fonts	
India	Assamese	as_IN.UTF-8	Lohit Bengali	fonts-bengali, scim-m17n, m17n-db-assamese
	Bengali	bn_IN.UTF-8	Lohit Bengali	fonts-bengali, scim-m17n, m17n-db-bengali
	Gujarati	gu_IN.UTF-8	Lohit Gujarati	fonts-gujarati, scim-m17n, m17n-db-gujarati
	Hindi	hi_IN.UTF-8	Lohit Hindi	fonts-hindi, scim-m17n, m17n-db-hindi
	Kannada	kn_IN.UTF-8	Lohit Kannada	fonts-kannada, scim-m17n, m17n-db-kannada
	Malayalam	ml_IN.UTF-8	Lohit Malayalam	fonts-malayalam, scim-m17n, m17n-db-malayalam
	Marathi	mr_IN.UTF-8	Lohit Hindi	fonts-hindi, scim-m17n, m17n-db-marathi
	Oriya	or_IN.UTF-8	Lohit Oriya	fonts-oriya, scim-m17n, m17n-db-oriya
	Punjabi	pa_IN.UTF-8	Lohit Punjabi	fonts-punjabi, scim-m17n, m17n-db-punjabi
	Tamil	ta_IN.UTF-8	Lohit Tamil	fonts-tamil, scim-m17n, m17n-db-tamil
Telugu	te_IN.UTF-8	Lohit Telugu	fonts-telugu, scim-m17n, m17n-db-telugu	

CHAPTER 2. 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 6.5. These changes include added or updated **procfs** entries, **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes.

reserved_blocks

This RW file contains a number of reserved blocks in the file system which are used in specific situations to avoid unexpected No space left on device (ENOSPC) errors or possible data loss.

proc/<pid>/comm and /proc/<pid>/task/<tid>/comm files

These files provide a method to access a task's comm value. It also allows for a task to set its own or one of its thread siblings' comm values. The comm value is limited in size compared to the cmdline value, so writing anything longer than the kernel's TASK_COMM_LEN macro (currently 16 chars) will result in a truncated comm value.

int_pln_enable

This parameter allows users to enable power limit notification interrupts.

nfsd.nfs4_disable_idmapping

The default value of this parameter is 0. When set to 1, NFSv4 server returns only numeric user IDs (UIDs) and group IDs (GIDs) to clients using AUTH_SYS mode, and will accept numeric UIDs and GIDs from such clients. This facilitates migration from NFS version 2 to NFS version 3.

PCI Subsystem Options

The following options for the **pci** kernel parameter can be used in Red Hat Enterprise Linux 6.5:

- **pcie_bus_tune_off**—disables PCIe maximum payload size (MPS) tuning and uses the BIOS-configured MPS default values.
- **pcie_bus_safe**—sets every device MPS to the largest value supported by all devices below the root complex.
- **pcie_bus_perf**—sets the device MPS to the largest allowable MPS based on its parent bus.
- **pcie_bus_peer2peer**— sets every device's MPS to 128B, which every device is guaranteed to support.

smbios_26_uuid

With this parameter, universally unique identifiers (UUIDs) are displayed in the System Management BIOS (SMBIOS) 2.6 format.

tsc_init_debug

With this parameter, additional information about the Time Stamp Counter (TSC) is displayed during system boot.

usbcore.usbfs_memory_mb

This option displays memory limit in MB for buffers allowed by USB device file system (usbfs).

tcp_limit_output_bytes

tcp_limit_output_bytes controls TCP Small Queue limit per TCP socket.

tcp_challenge_ack_limit

tcp_challenge_ack_limit limits the number of challenge acknowledgements sent per second, as recommended in RFC 5961 (Improving TCP's Robustness to Blind In-Window Attacks).

accept_ra

The **accept_ra** boolean allows for accepting router discovery messages (also known as router advertisements).

cookie_hmac_alg

cookie_hmac_alg is used to select the keyed-hash message authentication code (HMAC) algorithm used when generating the cookie value sent by a listening SCTP socket to a connecting client in the INIT-ACK chunk. Valid values are:

- md5
- sha1
- none

nf_conntrack_acct

The **nf_conntrack_acct** boolean enables connection tracking flow accounting.

nf_conntrack_buckets

nf_conntrack_buckets determines the size of a hash table. If it is not specified as parameter during module loading, the default size is calculated by dividing total memory by 16384 to determine the number of buckets but the hash table will never have fewer than 32 nor more than 16384 buckets.

nf_conntrack_checksum

This parameter is used to verify the checksums of incoming packets. Packets with invalid checksums are in INVALID state. If this is enabled, such packets will not be considered for connection tracking.

nf_conntrack_events_retry_timeout

This option is only relevant when "reliable connection tracking events" are used. Normally, ctnetlink is "lossy", that is, events are normally dropped when userspace listeners cannot keep up. Userspace can request "reliable event mode". When this mode is active, the connection tracking will only be destroyed after the event was delivered. If event delivery fails, the kernel periodically re-tries to send the event to userspace. The default value 15 is the maximum interval the kernel should use when re-trying to deliver the destroy event. A higher number means there will be fewer delivery retries and it will take longer for a backlog to be processed.

merge_across_nodes

The **merge_across_nodes** parameter specifies if pages from different NUMA nodes can be merged. When set to 0, Kernel SamePage Merging (KSM) merges only pages which physically reside in the memory area of the same NUMA node. 1 is the default value and merging across nodes

is performed as in earlier releases.

CHAPTER 3. DEVICE DRIVERS

This chapter provides a comprehensive listing of all device drivers which were updated in Red Hat Enterprise Linux 6.5.

Storage Drivers

- The Emulex **be2iscsi** driver has been upgraded to the latest upstream version.
- The **megaraid_sas** driver has been upgraded to version 6.600.18.00.
- The pm8001/pm80xx driver has been added in Red Hat Enterprise Linux 6.5 to add support for PMC-Sierra Adaptec Series 6H and 7H SAS/SATA HBA cards as well as PMC Sierra 8081, 8088, and 8089 chip based SAS/SATA controllers.
- 12Gbps SAS devices from LSI are now supported in Red Hat Enterprise Linux.
- The Brocade **BFA** driver has been updated to version 3.2.21.1.
- The **NVMe** driver has been added to Red Hat Enterprise Linux 6.
- Starting in Red Hat Enterprise Linux 6.4, iSCSI and FCoE boot on Broadcom devices is fully supported. These two features are provided by the bnx2i and bnx2fc Broadcom drivers.

Network Drivers

- The Virtual Extensible LAN, **vxlan**, driver has been updated.
- Support for Single Root I/O virtualization (SR-IOV) has been added to the **qlcnic** driver as a Technology Preview.
- The Brocade **BNA** driver has been updated to version 3.1.2.1.
- The **ixgbev** driver has been updated to the latest upstream version.
- The **igbvf** driver has been updated to the latest upstream version.
- The **bnx2x** driver has been updated to version 1.78.17-0.
- The Emulex **be2net** driver has been updated to version 4.6.x.
- The **qlcnic** driver has been updated to add support for the QLogic 83XX CNA adapter.
- The **e1000e** driver has been updated to the latest upstream version.
- The **tg3** driver has been updated to include various bug fixes and new features, including hardware PTP support.
- The **sfc** driver has been upgraded to upstream version 3.2 and includes hardware accelerated receive flow steering (RFS).
- The **igb** driver has been updated to version 4.1.2 to include software time stamping support.
- The **qlge** driver has been updated to version 1.00.00.32.

Miscellaneous Drivers

- The **hpilo** driver has been upgraded to the latest upstream version.

CHAPTER 4. TECHNOLOGY PREVIEWS

This chapter provides a list of all available Technology Previews in Red Hat Enterprise Linux 6.5.

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

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

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

4.1. STORAGE AND FILE SYSTEMS

Cross Realm Kerberos Trust Functionality for samba4 Libraries

The Cross Realm Kerberos Trust functionality provided by Identity Management, which relies on the capabilities of the samba4 client library, is included as a Technology Preview starting with Red Hat Enterprise Linux 6.4. This functionality uses the libndr-nbt library to prepare Connection-less Lightweight Directory Access Protocol (CLDAP) messages.

Package: samba-3.6.9-164

System Information Gatherer and Reporter (SIGAR)

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

Package: sigar-1.6.5-0.4.git58097d9

DIF/DIX support

DIF/DIX, is a new addition to the SCSI Standard and a Technology Preview in Red Hat Enterprise Linux 6. DIF/DIX increases the size of the commonly used 512-byte disk block from 512 to 520 bytes, adding the Data Integrity Field (DIF). The DIF stores a checksum value for the data block that is calculated by the Host Bus Adapter (HBA) when a write occurs. The storage device then confirms the checksum on receive, and stores both the data and the checksum. Conversely, when a read occurs, the checksum can be checked by the storage device, and by the receiving HBA.

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

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

Package: kernel-2.6.32-431

Btrfs, BZ#614121

Btrfs is under development as a file system capable of addressing and managing more files, larger files, and larger volumes than the ext2, ext3, and ext4 file systems. Btrfs is designed to make the file system tolerant of errors, and to facilitate the detection and repair of errors when they occur. It uses checksums to ensure the validity of data and metadata, and maintains snapshots of the file system that can be used for backup or repair. The Btrfs Technology Preview is only available on AMD64 and Intel 64 architectures.



WARNING

Red Hat Enterprise Linux 6 includes Btrfs as a technology preview to allow you to experiment with this file system. You should not choose Btrfs for partitions that will contain valuable data or that are essential for the operation of important systems.

Package: `btrfs-progs-0.20-0.2.git91d9eec`

LVM Application Programming Interface (API)

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

Package: `lvm2-2.02.100-8`

FS-Cache

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

Package: `cachefilesd-0.10.2-1`

eCryptfs File System

eCryptfs is a stacked, cryptographic file system. It is transparent to the underlying file system and provides per-file granularity. eCryptfs is provided as a Technology Preview in Red Hat Enterprise Linux 6.

Package: `ecryptfs-utils-82-6`

4.2. NETWORKING

Mellanox SR-IOV Support

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

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

Package: kernel-2.6.32-335

Open multicast ping (Omping), BZ#657370

Open Multicast Ping (Omping) is a tool to test the IP multicast functionality, primarily in the local network. This utility allows users to test IP multicast functionality and assists in the diagnosing if an issues is in the network configuration or elsewhere (that is, a bug). In Red Hat Enterprise Linux 6 Omping is provided as a Technology Preview.

Package: omping-0.0.4-1

QFQ queuing discipline

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

Package: kernel-2.6.32-431

vios-proxy, BZ#721119

vios-proxy is a stream-socket proxy for providing connectivity between a client on a virtual guest and a server on a Hypervisor host. Communication occurs over virtio-serial links.

Package: vios-proxy-0.2-1

4.3. CLUSTERING AND HIGH AVAILABILITY

luci support for fence_sanlock

The **luci** tool now supports the sanlock fence agent as a Technology Preview. The agent is available in the luci's list of agents.

Package: luci-0.26.0-48

Recovering a node via a hardware watchdog device

New fence_sanlock agent and checkquorum.wdmd, included in Red Hat Enterprise Linux 6.4 as a Technology Preview, provide new mechanisms to trigger the recovery of a node via a hardware watchdog device. Tutorials on how to enable this Technology Preview will be available at <https://fedorahosted.org/cluster/wiki/HomePage>

Note that SELinux in enforcing mode is currently not supported.

Package: cluster-3.0.12.1-59

keepalived

The keepalived package has been included as a Technology Preview, starting with Red Hat Enterprise Linux 6.4. The keepalived package provides simple and robust facilities for load-balancing and high-availability. The load-balancing framework relies on the well-know and widely used Linux Virtual Server kernel module providing Layer4 network load-balancing. The **keepalived** daemon implements a set of health checkers to load-balanced server pools according to their state. The keepalived daemon also implements the Virtual Router Redundancy Protocol (VRRP), allowing router or director failover to achieve high availability.

Package: keepalived-1.2.7-3

HAProxy

HAProxy is a stand-alone, layer-7, high-performance network load balancer for TCP and HTTP-based applications which can perform various types of scheduling based on the content of the HTTP requests. The haproxy package is included as a Technology Preview, starting with Red Hat Enterprise Linux 6.4.

Package: haproxy-1.4.24-2

4.4. AUTHENTICATION

Simultaneous maintaining of TGTs for multiple KDCs

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

Package: sssd-1.9.2-129

4.5. SECURITY

TPM

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

Packages: trousers-0.3.4-4, tpm-tools-1.3.4-2

4.6. DEVICES

mpt2sas lockless mode

The **mpt2sas** driver is fully supported. However, when used in the lockless mode, the driver is a Technology Preview.

Package: kernel-2.6.32-431

4.7. KERNEL

Kernel Media support

The following features are presented as Technology Previews:

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

Package: kernel-2.6.32-431

Linux (NameSpace) Container [LXC]

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

Packages: libvirt-0.9.10-21, virt-manager-0.9.0-14

Diagnostic pulse for the fence_ipmilan agent, [BZ#655764](#)

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

Package: fence-agents-3.1.5-35

CHAPTER 5. DEPRECATED FUNCTIONALITY

virtio-win component, BZ#1001981

The VirtIO SCSI driver has been removed from the virtio-win package and is no longer supported on Microsoft Windows Server 2003 platform.

qemu - kvm component

The qemu-guest-agent-win32 package is no longer shipped as part of the qemu-kvm package. The Windows guest agent is now delivered in the Supplementary channel together with other Windows components, for example, virtio-win drivers.

fence - agents component

Prior to Red Hat Enterprise Linux 6.5 release, the Red Hat Enterprise Linux High Availability Add-On was considered fully supported on certain VMware ESXi/vCenter versions in combination with the fence_scsi fence agent. Due to limitations in these VMware platforms in the area of SCSI-3 persistent reservations, the **fence_scsi** fencing agent is no longer supported on any version of the Red Hat Enterprise Linux High Availability Add-On in VMware virtual machines, except when using iSCSI-based storage. See the Virtualization Support Matrix for High Availability for full details on supported combinations:

<https://access.redhat.com/site/articles/29440>

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

systemtap component

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

matahari component

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

distribution component

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

- mingw-gcc
- mingw-boost
- mingw32-qpid-cpp
- python-qmf
- python-qpid
- qpid-cpp

- qpid-qmf
- qpid-tests
- qpid-tools
- ruby-qpid
- saslwrapper

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

fence-virt component

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

openscap component

The openscap-perl subpackage has been removed from openscap.

CHAPTER 6. KNOWN ISSUES

6.1. INSTALLATION

dracut component

For Fibre Channel over Ethernet (FCoE) from SAN on Dell systems which enable `biosdevname=1` by default, udev typically renames all network interfaces to their biosdevname naming convention during boot. However, a bug in udev prevents the FCoE boot interface from being renamed. This can result in occasional shutdown stalls. In order to install and shut down the system correctly, it is recommended to use the `biosdevname=0` installation parameter to avoid biosdevname naming in this case.

dracut component

For iSCSI boot from SAN on Dell systems which enable setting `biosdevname=1` by default, the installation completes successfully, but the system will not be able to mount the rootfs partition after reboot. This is because of a bug in Dracut where the boot network interface is not brought up if biosdevname naming is used. In order to install and reboot the system successfully in this case, use the `biosdevname=0` installation parameter to avoid biosdevname naming.

anaconda component

Setting the qla4xxx parameter `ql4xdisablesysfsboot` to `1` may cause boot from SAN failures.

anaconda component

To automatically create an appropriate partition table on disks that are uninitialized or contain unrecognized formatting, use the `zerombr` kickstart command. The `--initlabel` option of the `cleartpart` command is not intended to serve this purpose.

anaconda component, BZ#676025

Users performing an upgrade using the Anaconda's text mode interface who do not have a boot loader already installed on the system, or who have a non-GRUB boot loader, need to select **Skip Boot Loader Configuration** during the installation process. Boot loader configuration will need to be completed manually after installation. This problem does not affect users running Anaconda in the graphical mode (graphical mode also includes VNC connectivity mode).

anaconda component

On s390x systems, you cannot use automatic partitioning and encryption. If you want to use storage encryption, you must perform custom partitioning. Do not place the `/boot` volume on an encrypted volume.

anaconda component

The order of device names assigned to USB attached storage devices is not guaranteed. Certain USB attached storage devices may take longer to initialize than others, which can result in the device receiving a different name than you expect (for example, `sdc` instead of `sda`).

During installation, verify the storage device size, name, and type when configuring partitions and file systems.

kernel component

Recent Red Hat Enterprise Linux 6 releases use a new naming scheme for network interfaces on

some machines. As a result, the installer may use different names during an upgrade in certain scenarios (typically **em1** is used instead of **eth0** on new Dell machines). However, the previously used network interface names are preserved on the system and the upgraded system will still use the previously used interfaces. This is not the case for Yum upgrades.

anaconda component

The **kdump default on** feature currently depends on Anaconda to insert the **crashkernel=** parameter to the kernel parameter list in the boot loader's configuration file.

firstaidkit component

The **firstaidkit-plugin-grub** package has been removed from Red Hat Enterprise Linux 6.2. As a consequence, in rare cases, the system upgrade operation may fail with unresolved dependencies if the plug-in has been installed in a previous version of Red Hat Enterprise Linux. To avoid this problem, the **firstaidkit-plugin-grub** package should be removed before upgrading the system. However, in most cases, the system upgrade completes as expected.

anaconda component, BZ#623261

In some circumstances, disks that contain a whole disk format (for example, an LVM Physical Volume populating a whole disk) are not cleared correctly using the **clearpart --initlabel** kickstart command. Adding the **--all** switch—as in **clearpart --initlabel --all**—ensures disks are cleared correctly.

anaconda component

When installing on the IBM System z architecture, if the installation is being performed over SSH, avoid resizing the terminal window containing the SSH session. If the terminal window is resized during the installation, the installer will exit and the installation will terminate.

yaboot component, BZ#613929

The kernel image provided on the CD/DVD is too large for Open Firmware. Consequently, on the POWER architecture, directly booting the kernel image over a network from the CD/DVD is not possible. Instead, use **yaboot** to boot from a network.

anaconda component

The Anaconda partition editing interface includes a button labeled **Resize**. This feature is intended for users wishing to shrink an existing file system and an underlying volume to make room for an installation of a new system. Users performing manual partitioning cannot use the **Resize** button to change sizes of partitions as they create them. If you determine a partition needs to be larger than you initially created it, you must delete the first one in the partitioning editor and create a new one with the larger size.

system-config-kickstart component

Channel IDs (read, write, data) for network devices are required for defining and configuring network devices on IBM S/390 systems. However, **system-config-kickstart**—the graphical user interface for generating a kickstart configuration—cannot define channel IDs for a network device. To work around this issue, manually edit the kickstart configuration that **system-config-kickstart** generates to include the desired network devices.

6.2. ENTITLEMENT

subscription-manager component

If multiple repositories are enabled, **subscription-manager** installs product certificates from all repositories instead of installing the product certificate only from the repository from which the RPM package was installed.

6.3. DEPLOYMENT

389-ds-base component, BZ#878111

The **ns-slapd** utility terminates unexpectedly if it cannot rename the **dirsrv-<instance>** log files in the **/var/log/** directory due to incorrect permissions on the directory.

cpuspeed component, BZ#626893

Some HP Proliant servers may report incorrect CPU frequency values in **/proc/cpuinfo** or **/sys/device/system/cpu/*/cpufreq**. This is due to the firmware manipulating the CPU frequency without providing any notification to the operating system. To avoid this ensure that the **HP Power Regulator** option in the BIOS is set to **OS Control**. An alternative available on more recent systems is to set **Collaborative Power Control** to **Enabled**.

releng component, BZ#644778

Some packages in the Optional repositories on RHN have multilib file conflicts. Consequently, these packages cannot have both the primary architecture (for example, x86_64) and secondary architecture (for example, i686) copies of the package installed on the same machine simultaneously. To work around this issue, install only one copy of the conflicting package.

grub component, BZ#695951

On certain UEFI-based systems, you may need to type **BOOTX64** rather than **bootx64** to boot the installer due to case sensitivity issues.

grub component, BZ#698708

When rebuilding the grub package on the x86_64 architecture, the glibc-static.i686 package must be used. Using the glibc-static.x86_64 package will not meet the build requirements.

6.4. VIRTUALIZATION

qemu-kvm component, BZ#1159613

If a **virtio** device is created where the number of vectors is set to a value higher than 32, the device behaves as if it was set to a zero value on Red Hat Enterprise Linux 6, but not on Enterprise Linux 7. The resulting vector setting mismatch causes a migration error if the number of vectors on any **virtio** device on either platform is set to 33 or higher. It is, therefore, not recommended to set the **vector** value to be greater than 32.

virtio-win component

When upgrading the **NetKVM** driver through the Windows Device Manager, the old registry values are not removed. As a consequence, for example, non-existent parameters may be available.

qemu-kvm component

When working with very large images (larger than 2TB) created with very small cluster sizes (for example, 512bytes), block I/O errors can occur due to timeouts in qemu. To prevent this problem from occurring, use the default cluster size of 64KiB or larger.

kernel component

On Microsoft Windows Server 2012 containing large dynamic VHDX (Hyper-V virtual hard disk) files and using the ext3 file system, a call trace can appear, and, consequently, it is not possible to shut down the guest. To work around this problem, use the ext4 file system or set a logical block size of 1MB when creating a VHDX file. Note that this can only be done by using Microsoft PowerShell as the Hyper-V manager does not expose the `-BlockSizeBytes` option which has the default value of 32MB. To create a dynamix VHDX file with an approximate size of 2.5TB and 1MB block size run:

```
New-VHD -Path .\MyDisk.vhdx -SizeBytes 5120MB -BlockSizeBytes 1MB -
Dynamic
```

libvirt component

The storage drivers do not support the `virsh vol-resize` command options `--allocate` and `--shrink`. Use of the `--shrink` option will result in the following error message:

```
error: invalid argument: storageVolumeResize: unsupported flags (0x4)
```

Use of the `--allocate` option will result in the following error message:

```
error: invalid argument: storageVolumeResize: unsupported flags (0x1)
```

Shrinking a volume's capacity is possible as long as the value provided on the command line is greater than the volume allocation value as seen with the `virsh vol-info` command. You can shrink an existing volume by name through the following sequence of steps:

1. Dump the XML of the larger volume into a file using the `vol-dumpxml` .
2. Edit the file to change the name, path, and capacity values, where the capacity must be greater than or equal to the allocation.
3. Create a temporary smaller volume using the `vol-create` with the edited XML file.
4. Back up and restore the larger volumes data using the `vol-download` and `vol-upload` commands to the smaller volume.
5. Use the `vol-delete` command to remove the larger volume.
6. Use the `vol-clone` command to restore the name from the larger volume.
7. Use the `vol-delete` command to remove the temporary volume.

In order to allocate more space on the volume, follow a similar sequence, but adjust the allocation to a larger value than the existing volume.

virtio-win component

It is not possible to downgrade a driver using the `Search for the best driver in these locations` option because the newer and installed driver will be selected as the "best" driver. If you want to force installation of a particular driver version, use the `Don't search` option and the `Have`

Disk button to select the folder of the older driver. This method will allow you to install an older driver on a system that already has a driver installed.

kernel component

There is a known issue with the Microsoft Hyper-V host. If a legacy network interface controller (NIC) is used on a multiple-CPU virtual machine, there is an interrupt problem in the emulated hardware when the IRQ balancing daemon is running. Call trace information is logged in the `/var/log/messages` file.

libvirt component, BZ#888635

Under certain circumstances, virtual machines try to boot from an incorrect device after a network boot failure. For more information, please refer to [this article](#) on Customer Portal.

numad component, BZ#872524

If **numad** is run on a system with a task that has very large resident memory ($\geq 50\%$ total system memory), then the numad-initiated NUMA page migrations for that task can cause swapping. The swapping can then induce long latencies for the system. An example is running a 256GB Microsoft Windows KVM Virtual Machine on a 512GB host. The Windows guest will fault in all pages on boot in order to zero them. On a four node system, **numad** will detect that a 256GB task can fit in a subset of two or three nodes, and then attempt to migrate it to that subset. Swapping can then occur and lead to latencies. These latencies may then cause the Windows guest to hang, as timing requirements are no longer met. Therefore, on a system with only one or two very large Windows machines, it is recommended to disable **numad**.

Note that this problem is specific to Windows 2012 guests that use more memory than exists in a single node. Windows 2012 guests appear to allocate memory more gradually than other Windows guest types, which triggers the issue. Other varieties of Windows guests do not seem to experience this problem. You can work around this problem by:

- limiting Windows 2012 guests to less memory than exists in a given node -- so on a typical 4 node system with even memory distribution, the guest would need to be less than the total amount of system memory divided by 4; or
- allowing the Windows 2012 guests to finish allocating all of its memory before allowing **numad** to run. **numad** will handle extremely huge Windows 2012 guests correctly after allowing a few minutes for the guest to finish allocating all of its memory.

grubby component, BZ#893390

When a Red Hat Enterprise Linux 6.4 guest updates the kernel and then the guest is turned off through Microsoft Hyper-V Manager, the guest fails to boot due to incomplete grub information. This is because the data is not synced properly to disk when the machine is turned off through Hyper-V Manager. To work around this problem, execute the **sync** command before turning the guest off.

kernel component

Using the mouse scroll wheel does not work on Red Hat Enterprise Linux 6.4 guests that run under certain version of Microsoft Hyper-V Manager. However, the scroll wheel works as expected when the **vncviewer** utility is used.

kernel component, BZ#874406

Microsoft Windows Server 2012 guests using the e1000 driver can become unresponsive consuming 100% CPU during boot or reboot.

kernel component

When a kernel panic is triggered on a Microsoft Hyper-V guest, the **kdump** utility does not capture the kernel error information; an error is only displayed on the command line. This is a host problem. Guest **kdump** works as expected on Microsoft Hyper-V 2012 R2 host.

qemu - kvm component, BZ#871265

AMD Opteron G1, G2 or G3 CPU models on **qemu-kvm** use the family and models values as follows: family=15 and model=6. If these values are larger than 20, the **lahfm_lm** CPU feature is ignored by Linux guests, even when the feature is enabled. To work around this problem, use a different CPU model, for example AMD Opteron G4.

qemu - kvm component, BZ#860929

KVM guests must not be allowed to update the host CPU microcode. KVM does not allow this, and instead always returns the same microcode revision or patch level value to the guest. If the guest tries to update the CPU microcode, it will fail and show an error message similar to:

```
CPU0: update failed (for patch_level=0x6000624)
```

To work around this, configure the guest to not install CPU microcode updates; for example, uninstall the `microcode_ctl` package Red Hat Enterprise Linux or Fedora guests.

virt - p2v component, BZ#816930

Converting a physical server running either Red Hat Enterprise Linux 4 or Red Hat Enterprise Linux 5 which has its file system root on an MD device is not supported. Converting such a guest results in a guest which fails to boot. Note that conversion of a Red Hat Enterprise Linux 6 server which has its root on an MD device is supported.

virt - p2v component, BZ#808820

When converting a physical host with a multipath storage, Virt-P2V presents all available paths for conversion. Only a single path must be selected. This must be a currently active path.

virtio-win component, BZ#615928

The balloon service on Windows 7 guests can only be started by the Administrator user.

libvirt component, BZ#622649

libvirt uses transient **iptables** rules for managing NAT or bridging to virtual machine guests. Any external command that reloads the **iptables** state (such as running **system-config-firewall**) will overwrite the entries needed by **libvirt**. Consequently, after running any command or tool that changes the state of **iptables**, guests may lose access to the network. To work around this issue, use the **service libvirt reload** command to restore **libvirt**'s additional **iptables** rules.

virtio-win component, BZ#612801

A Windows virtual machine must be restarted after the installation of the kernel Windows driver framework. If the virtual machine is not restarted, it may crash when a memory balloon operation is performed.

qemu - kvm component, BZ#720597

Installation of Windows 7 Ultimate x86 (32-bit) Service Pack 1 on a guest with more than 4GB of RAM and more than one CPU from a DVD medium can lead to the system being unresponsive and,

consequently, to a crash during the final steps of the installation process. To work around this issue, use the Windows Update utility to install the Service Pack.

qemu - kvm component, BZ#612788

A dual function Intel 82576 Gigabit Ethernet Controller interface (codename: Kawela, PCI Vendor/Device ID: 8086:10c9) cannot have both physical functions (PF's) device-assigned to a Windows 2008 guest. Either physical function can be device assigned to a Windows 2008 guest (PCI function 0 or function 1), but not both.

virt -v2v component, BZ#618091

The **virt-v2v** utility is able to convert guests running on an ESX server. However, if an ESX guest has a disk with a snapshot, the snapshot must be on the same datastore as the underlying disk storage. If the snapshot and the underlying storage are on different datastores, **virt-v2v** will report a 404 error while trying to retrieve the storage.

virt -v2v component, BZ#678232

The VMware Tools application on Microsoft Windows is unable to disable itself when it detects that it is no longer running on a VMware platform. Consequently, converting a Microsoft Windows guest from VMware ESX, which has VMware Tools installed, will result in errors. These errors usually manifest as error messages on start-up, and a "Stop Error" (also known as a BSOD) when shutting down the guest. To work around this issue, uninstall VMware Tools on Microsoft Windows guests prior to conversion.

libguestfs component

The libguestfs packages do not support remote access to disks over the network in Red Hat Enterprise Linux 6. Consequently, the **virt-sysprep** tool as well as other tools do not work with remote disks. Users who need to access disks remotely with tools such as **virt-sysprep** are advised to upgrade to Red Hat Enterprise Linux 7.

6.5. STORAGE AND FILE SYSTEMS

lvm2 component, BZ#1024347

An event is generated for any device that is being watched for changes by means of a special WATCH udev rule. This udev rule is also used for logical volumes and it causes the `/dev/` directory to be up-to-date with any data written to the logical volume (mainly the symlinks that are based on metadata, like the content of the `/dev/disk` directory). The event is generated each time the device is closed after being open for writing.

```
device-mapper: remove ioctl on failed: Device or resource busy
```

This is caused by the LVM command and udev interaction where the original logical volume is open for writing and then part of the logical volume is zeroed so it is prepared for thin pool use. Then the logical volume is closed, which triggers the WATCH rule. Then LVM tries to remove the original volume while it can still be opened by udev. This causes the error message to appear. LVM tries to remove the logical volume a few times before exiting with an **lvconvert** failure. Normally, udev should process the logical volume quickly and LVM should continue retrying to remove the logical volume. Normally, users can just ignore this error message; the logical volume is processed correctly on next retry. If the number of retries is not sufficient, then **lvconvert** can fail as a result. If this is the case, users are encouraged to comment out the **OPTIONS+="watch"** line in the `/lib/udev/rules.d/13-dm-disk.rules` file. This will cause the WATCH rule for LVM volumes to be disabled. However, this may cause the `/dev/` content to be out-of-sync with actual metadata

state stored on the logical volume. If LVM needs to retry the logical volume removal because it is being opened in parallel, most notably by `udev` as described before, it issues an error message "remove ioctl failed: Device or resource busy". If this is the case, the removal is retried several times before `lvconvert` fails completely.

device-mapper-persistent-data component, BZ#960284

Tools provided by the `device-mapper-persistent-data` package fail to operate on 4K hard-sectored metadata devices.

anaconda component

In UEFI mode, when creating a partition for software RAID, `anaconda` can be unable to allocate the `/boot/efi` mount point to the software RAID partition and fails with the "have not created /boot/efi" message in such a scenario.

kernel component, BZ#918647

Thin provisioning uses reference counts to indicate that data is shared between a thin volume and snapshots of the thin volume. There is a known issue with the way reference counts are managed in the case when a discard is issued to a thin volume that has snapshots. Creating snapshots of a thin volume and then issuing discards to the thin volume can therefore result in data loss in the snapshot volumes. Users are strongly encouraged to disable discard support on the thin-pool for the time being. To do so using `lvm2` while the pool is offline, use the `lvchange --discard ignore <pool>` command. Any discards that might be issued to thin volumes will be ignored.

kernel component

Storage that reports a `discard_granularity` that is not a power of two will cause the kernel to improperly issue discard requests to the underlying storage. This results in I/O errors associated with the failed discard requests. To work around the problem, if possible, do not upgrade to newer vendor storage firmware that reports `discard_granularity` that is not a power of two.

parted component

Users might be unable to access a partition created by `parted`. To work around this problem, reboot the machine.

lvm2 component, BZ#852812

When filling a thin pool to 100% by writing to thin volume device, access to all thin volumes using this thin pool can be blocked. To prevent this, try not to overfill the pool. If the pool is overfilled and this error occurs, extend the thin pool with new space to continue using the pool.

dracut component

The Qlogic QLA2xxx driver can miss some paths after booting from Storage Area Network (SAN). To workaroud this problem, run the following commands:

```
echo "options qla2xxx ql2xasynclgin=0" > /etc/modprobe.d/qla2xxx.conf
mkinitrd /boot/initramfs-`uname -r`.img `uname -r` --force
```

lvm2 component, BZ#903411

Activating a logical volume can fail if the `--thinpool` and `--discards` options are specified on logical-volume creation. To work around this problem, manually deactivate all thin volumes related to the changed thin pool prior to running the `lvchange` command.

kernel component

Unloading the **nfs** module can cause the system to terminate unexpectedly if the **fsx** utility was ran with NFSv4.1 before.

device-mapper-multipath component

When the **multipathd** service is not running, failed devices will not be restored. However, the **multipath** command gives no indication that **multipathd** is not running. Users can unknowingly set up **multipath** devices without starting the **multipathd** service, keeping failed paths from automatically getting restored. Make sure to start multipathing by

- either running:

```
~]# mpathconf --enable
~]# service multipathd start
```

- or:

```
~]# chkconfig multipathd on
~]# service multipathd start
```

multipathd will automatically start on boot, and **multipath** devices will automatically restore failed paths.

lvm2 component, BZ#837603

When the administrator disables use of the **lvmetad** daemon in the **lvm.conf** file, but the daemon is still running, the cached metadata are remembered until the daemon is restarted. However, if the **use_lvmetad** parameter in **lvm.conf** is reset to **1** without an intervening **lvmetad** restart, the cached metadata can be incorrect. Consequently, VG metadata can be overwritten with previous versions. To work around this problem, stop the **lvmetad** daemon manually when disabling **use_lvmetad** in **lvm.conf**. The daemon can only be restarted after **use_lvmetad** has been set to **1**. To recover from an out-of-sync **lvmetad** cache, execute the **pvscan --cache** command or restart **lvmetad**. To restore metadata to correct versions, use **vgcfrestore** with a corresponding file in **/etc/lvm/archive**.

lvm2 component, BZ#563927

Due to the limitations of the LVM 'mirror' segment type, it is possible to encounter a deadlock situation when snapshots are created of mirrors. The deadlock can occur if snapshot changes (e.g. creation, resizing or removing) happen at the same time as a mirror device failure. In this case, the mirror blocks I/O until LVM can respond to the failure, but the snapshot is holding the LVM lock while trying to read the mirror.

If the user wishes to use mirroring and take snapshots of those mirrors, then it is recommended to use the 'raid1' segment type for the mirrored logical volume instead. This can be done by adding the additional arguments '--type raid1' to the command that creates the mirrored logical volume, as follows:

```
~]$ lvcreate --type raid1 -m 1 -L 1G -n my_mirror my_vg
```

kernel component, BZ#606260

The NFSv4 server in Red Hat Enterprise Linux 6 currently allows clients to mount using UDP and advertises NFSv4 over UDP with **rpcbind**. However, this configuration is not supported by Red Hat and violates the RFC 3530 standard.

lvm2 component

The **pvmove** command cannot currently be used to move mirror devices. However, it is possible to move mirror devices by issuing a sequence of two commands. For mirror images, add a new image on the destination PV and then remove the mirror image on the source PV:

```
~]$ lvconvert -m +1 <vg/lv> <new PV>
~]$ lvconvert -m -1 <vg/lv> <old PV>
```

Mirror logs can be handled in a similar fashion:

```
~]$ lvconvert --mirrorlog core <vg/lv>
~]$ lvconvert --mirrorlog disk <vg/lv> <new PV>
```

or

```
~]$ lvconvert --mirrorlog mirrored <vg/lv> <new PV>
~]$ lvconvert --mirrorlog disk <vg/lv> <old PV>
```

6.6. NETWORKING

kernel component

In cluster environment, the multicast traffic from the guest to a host can be unreliable. To work around this problem, enable `multicast_querier` for the bridge. The setting is located in the `/sys/class/net/<bridge_name>/bridge/multicast_querier` file. Note that if the setting is not available, the problem should not occur.

kernel component

A missing part of the **bcma** driver causes the **brcmsmac** driver not to load automatically when the **bcma** driver scans the for devices. This causes the kernel not to load the **brcmsmac** module automatically on boot. Symptoms can be confirmed by running the `lspci -v` command for the device and noting the driver to be **bmca**, not **brcmsmac**. To load the driver manually, run `modprobe brcmsmac` on the command line.

389-ds-base component

Under certain conditions, when the server is processing multiple outgoing replication or windows sync agreements using the TLS or SSL protocol, and processing incoming client requests that use TLS or SSL and Simple Paged Results, the server becomes unresponsive to new incoming client requests. The **dirsrv** service will stop responding to new incoming client requests. A restart of the **dirsrv** service is required to restore service.

kernel component, BZ#1003475

When some Fibre Channel over Ethernet (FCoE) switch ports connected to the bfa host bus adapter go offline and then return in the online state, the bfa port may not re-establish the connection with the switch. This is due to a failure of the bfa driver's retry logic when interacting with certain switches. To work around this problem, reset the bfa link. This can be done either by running:

-

```
]# echo 1 > /sys/class/fc_host/host/issue_lip
```

or by running:

```
]# modprobe -r bfa && modprobe bfa
```

anaconda component, BZ#984129

For HP systems running in HP FlexFabric mode, the designated iSCSI function can only be used for iSCSI offload related operations and will not be able to perform any other Layer 2 networking tasks, for example, DHCP. In the case of iSCSI boot from SAN, the same SAN MAC address is exposed to both the corresponding **ifconfig** record and the iSCSI Boot Firmware Table (iBFT), therefore, Anaconda will skip the network selection prompt and will attempt to acquire the IP address as specified by iBFT. If DHCP is desired, Anaconda will attempt to acquire DHCP using this iSCSI function, which will fail and Anaconda will then try to acquire DHCP indefinitely. To work around this problem, if DHCP is desired, the user must use the **asknetwork** installation parameter and provide a "dummy" static IP address to the corresponding network interface of the iSCSI function. This prevents Anaconda from entering an infinite loop and allows it to request the iSCSI offload function to perform DHCP acquisition instead.

iscsi-initiator-utils component, BZ#825185

If the corresponding network interface has not been brought up by **dracut** or the tools from the **iscsi-initiator-utils** package, this prevents the correct MAC address from matching the offload interface, and host bus adapter (HBA) mode will not work without manual intervention to bring the corresponding network interface up. To work around this problem, the user must select the corresponding Layer 2 network interface when **anaconda** prompts the user to choose "which network interface to install through". This will inherently bring up the offload interface for the installation.

kernel component

When an **igb** link is up, the following **ethtool** fields display incorrect values as follows:

- *Supported ports: []* - for example, an empty bracket can be displayed.
- *Supported pause frame use: No* - however, pause frame is supported.
- *Supports auto-negotiation: No* - auto-negotiation is supported.
- *Advertised pause frame use: No* - advertised pause frame is turned on.
- *Advertised auto-negotiation: No* - advertised auto-negotiation is turned on.
- *Speed: Unknown!* - the speed is known and can be verified using the **dmesg** tool.

linuxptp component

End-to-End (E2E) slaves that communicated with an E2E master once can synchronize to Peer-to-Peer (P2P) masters and vice versa. The slaves cannot update their path delay value because E2E ports reject peer delay requests from P2P ports. However, E2E ports accept SYNC messages from P2P ports and the slaves keep updating clock frequency based on undesired offset values that are calculated by using the old path delay value. Therefore, a time gap will occur if the master port is started with an incorrect delay mechanism. The "delay request on P2P" or "pdelay_req on E2E port" message can appear. To work around these problems, use a single delay mechanism for one PTP communication path. Also, because E2E and P2P mismatch can trigger a time gap of slave clock, pay attention to the configuration when starting or restarting a node on a running domain.

samba4 component, BZ#878168

If configured, the Active Directory (AD) DNS server returns IPv4 and IPv6 addresses of an AD server. If the FreeIPA server cannot connect to the AD server with an IPv6 address, running the **ipa trust-add** command will fail even if it would be possible to use IPv4. To work around this problem, add the IPv4 address of the AD server to the **/etc/hosts** file. In this case, the FreeIPA server will use only the IPv4 address and executing **ipa trust-add** will be successful.

kernel component

Destroying the root port before any NPIV ports can cause unexpected system behavior, including a full system crash. Note that one instance where the root port is destroyed before the NPIV ports is when the system is shut down. To work around this problem, destroy NPIV ports before destroying the root port that the NPIV ports were created on. This means that for each created NPIV port, the user should write to the **sysfs vport_delete** interface to delete that NPIV port. This should be done before the root port is destroyed. Users are advised to script the NPIV port deletion and configure the system such that the script is executed before the **fcoe** service is stopped, in the shutdown sequence.

kernel component

A Linux LIO FCoE target causes the **bfa** driver to reset all FCoE targets which might lead to data corruption on LUN. To avoid these problems, do not use the **bfa** driver with a Linux FCoE target.

NetworkManager component, BZ#896198

A **GATEWAY** setting in the **/etc/sysconfig/network** file causes **NetworkManager** to assign that gateway to all interfaces with static IP addresses, even if their configuration did not specify a gateway or specified a different gateway. Interfaces have the incorrect gateway information and the wrong interface may have the default route. Instead of using **GATEWAY** in **/etc/sysconfig/network** to specify which interface receives the default route, set **DEFROUTE=no** in each **ifcfg** file that should *not* have the default route. Any interface connected using configuration from an **ifcfg** file containing **DEFROUTE=no** will never receive the default route.

kernel component

Typically, on platforms with no Intelligent Platform Management Interface (IPMI) hardware the user can see the following message on the boot console and in **dmesg** log:

```
Could not set up I/O space
```

This message can be safely ignored, unless the system really does have IPMI hardware. In that case, the message indicates that the IPMI hardware could not be initialized. In order to support Advanced Configuration and Power Interface (ACPI) opregion access to IPMI functionality early in the boot, the IPMI driver has been statically linked with the kernel image. This means that the IPMI driver is "loaded" whether or not there is any hardware. The IPMI driver will try to initialize the IPMI hardware, but if there is no IPMI hardware present on the booting platform, the driver will print error messages on the console and in the **dmesg** log. Some of these error messages do not identify themselves as having been issued by the IPMI driver, so they can appear to be serious, when they are harmless.

kernel component

Shutting down the **fcoe-target** service while the Fibre Channel over Ethernet (FCoE) can lead to a kernel crash. Please minimize FCoE traffic before stopping or restarting this service.

fcoe-utils component

After an ixgbe Fibre Channel over Ethernet (FCoE) session is created, server reboot can cause some or all of the FCoE sessions to not be created automatically. To work around this problem, follow the following steps (assuming that *eth0* is the missing NIC for the FCoE session):

```
ifconfig eth0 down
ifconfig eth0 up
sleep 5
dcbtool sc eth0 dcb on
sleep 5
dcbtool sc eth0 pfc e:1 a:1 w:1
dcbtool sc eth0 app:fcoe e:1 a:1 w:1
service fcoe restart
```

libibverbs component

The InfiniBand UD transport test utility could become unresponsive when the **ibv_ud_pingpong** command was used with a packet size of 2048 or greater. UD is limited to no more than the smallest MTU of any point in the path between point A and B, which is between 0 and 4096 given that the largest MTU supported (but not the smallest nor required) is 4096. If the underlying Ethernet is jumbo frame capable, and with a 4096 IB MTU on an RoCE device, the max packet size that can be used with UD is 4012 bytes.

bind-dyndb-ldap component

IPA creates a new DNS zone in two separate steps. When the new zone is created, it is invalid for a short period of time. **A/AAAA** records for the name server belonging to the new zone are created after this delay. Sometimes, **BIND** attempts to load this invalid zone and fails. In such a case, reload **BIND** by running either **rndc reload** or **service named restart**.

selinux-policy component

SELinux can prevent the **nmbd** service from writing into the **/var/**, which breaks NetBIOS name resolution and leads to SELinux AVC denials.

kernel component

The latest version of the **sfc** NIC driver causes lower UDP and TX performance with large amounts of fragmented UDP packets. This problem can be avoided by setting a constant interrupt moderation period (not adaptive moderation) on both sides, sending and receiving.

kernel component

Some network interface cards (NICs) may not get an IPv4 address assigned after the system is rebooted. To work around this issue, add the following line to the **/etc/sysconfig/network-scripts/ifcfg-*interface*** file:

```
LINKDELAY=10
```

NetworkManager component, BZ#758076

If a Certificate Authority (CA) certificate is not selected when configuring an 802.1x or WPA-Enterprise connection, a dialog appears indicating that a missing CA certificate is a security risk. This dialog presents two options: ignore the missing CA certificate and proceed with the insecure connection, or choose a CA certificate. If the user elects to choose a CA certificate, this dialog disappears and the user may select the CA certificate in the original configuration dialog.

samba component

Current Samba versions shipped with Red Hat Enterprise Linux 6 are not able to fully control the user and group database when using the **ldapsam_compat** back end. This back end was never designed to run a production LDAP and Samba environment for a long period of time. The **ldapsam_compat** back end was created as a tool to ease migration from historical Samba releases (version 2.2.x) to Samba version 3 and greater using the new **ldapsam** back end and the new LDAP schema. The **ldapsam_compat** back end lack various important LDAP attributes and object classes in order to fully provide full user and group management. In particular, it cannot allocate user and group IDs. In the [Red Hat Enterprise Linux Reference Guide](#), it is pointed out that this back end is likely to be deprecated in future releases. Refer to Samba's [documentation](#) for instructions on how to migrate existing setups to the new LDAP schema.

When you are not able to upgrade to the new LDAP schema (though upgrading is strongly recommended and is the preferred solution), you may work around this issue by keeping a dedicated machine running an older version of Samba (v2.2.x) for the purpose of user account management. Alternatively, you can create user accounts with standard LDIF files. The important part is the assignment of user and group IDs. In that case, the old Samba 2.2 algorithmic mapping from Windows RIDs to Unix IDs is the following: $user\ RID = UID * 2 + 1000$ while for groups it is: $group\ RID = GID * 2 + 1001$. With these workarounds, users can continue using the **ldapsam_compat** back end with their existing LDAP setup even when all the above restrictions apply.

kernel component

Because Red Hat Enterprise Linux 6 defaults to using Strict Reverse Path filtering, packets are dropped by default when the route for outbound traffic differs from the route of incoming traffic. This is in line with current recommended practice in RFC3704. For more information about this issue please refer to [/usr/share/doc/kernel-doc-<version>/Documentation/networking/ip-sysctl.txt](#) and <https://access.redhat.com/site/solutions/53031>.

6.7. SECURITY

kernel component

When stopping the **ipsec** daemon, error messages about modules being in use can occur.

openssl component, BZ#1022002

The external Advanced Encryption Standard (AES) New Instructions (AES-NI) engine is no longer available in openssl; the engine is now built-in and therefore no longer needs to be manually enabled.

6.8. CLUSTERING

corosync component

The redundant ring feature of corosync is not fully supported in combination with InfiniBand or Distributed Lock Manager (DLM). A double ring failure can cause both rings to break at the same time on different nodes. In addition, DLM is not functional if ring0 is down.

lvm2 component, BZ#814779

Clustered environment is not supported by **lvm2** at the moment. If `global/use_lvm2=1` is used together with `global/locking_type=3` configuration setting (clustered locking), the `use_lvm2` setting is automatically overridden to `0` and **lvm2** is not used in this case at all. Also, the following warning message is displayed:

WARNING: configuration setting `use_lvmetad` overridden to 0 due to `locking_type` 3. Clustered environment not supported by `lvmetad` yet.

luci component, BZ#615898

`luci` will not function with Red Hat Enterprise Linux 5 clusters unless each cluster node has `ricci` version 0.12.2-14.

6.9. AUTHENTICATION

ipa component, BZ#1024744

OpenLDAP and 389 Directory Server treat the grace logins differently. 389 Directory Server treats them as "number of grace logins left" while OpenLDAP treats them as "number of grace logins used". Currently the SSSD only handles the semantics used by 389 Directory server. As a result, when using OpenLDAP server, the grace password warning might be incorrect.

ipa component, BZ#1024959

The Identity Management server does not write the initial user password correctly to password history. As a consequence, when a new Identity Management user is created and a password is generated for him, the first time that user changes the password, the value of the first password is disregarded when the password policy plug-in checks the password history. This means that user can "change" the initial password to the same value as the previous one, with no regards to the configured password history. Password history is applied correctly to all subsequent password changes.

ipa component, BZ#1009102

When an Identity Management server installed on Red Hat Enterprise Linux 6.2 is updated to the version provided by Red Hat Enterprise Linux 6.4 or 6.5, the new `pbac` permission "Write DNS Configuration" is created without any of the required object classes. Consequently, the permission may not show up on the Identity Management Web UI permission page or when the `--sizelimit` parameter is used for the CLI `permission-find` command. The permission is still accessible using the command line when the `--sizelimit` option is not specified. To work around this problem, run the following command on the server to trigger the DNS permission update process again and fix the list of permission object classes:

```
]# ipa-ldap-updater --ldapi /usr/share/ipa/updates/40-dns.update
```

This problem can also be avoided when a Red Hat Enterprise Linux 6.4 or 6.5 replica is installed or when an Identity Management server is reinstalled or upgraded.

ipa component, BZ#1015481

Identity Management administration framework API contains two checks to verify that a request on its API can be passed further:

1. A check to see if the client API version is not higher than the server API version. If it is, the request is rejected.
2. A check to see if the client API request does not use an attribute or a parameter unknown to the server. If it does, the request is rejected.

However, the Identity Management server performs the checks in an incorrect order: first, the attribute and parameter check is done and after that, the API version check is done. As a consequence, when

a new client (for example, Red Hat Enterprise Linux 6.5) runs the **ipa** administration tool against a server with an earlier operating system (for example, Red Hat Enterprise Linux 6.4), the command returns a confusing error message; for example, instead of stating API compatibility, **ipa** outputs the following message:

```
]$ ipa user-show admin
ipa: ERROR: Unknown option: no_members
```

ipa component, BZ#1016042

The **ipa-replica-manage** tool contains a bug in the **re-initialize** command causing the MemberOf task to fail with an error under certain circumstances. When the **ipa-replica-manage re-initialize** command is run for a Windows Synchronization (WinSync) replication agreement, it succeeds in the re-initialization part, but fails during execution of the MemberOf task which is run after the re-initialization part. The following error is returned:

```
Update succeeded
Can't contact LDAP server
```

However, the error is harmless as running the MemberOf task is not required in this case.

sssd component, BZ#995737

SSSD fails if the entryUSN attribute of sudo rules is empty. As a result, processing of sudo rules stops instead of proceeding. To work around this problem, if the server contains any other USN-like attribute, the user can set the attribute in the configuration file using:

```
ldap_rootdse_last_usn = attr_name
ldap_entry_usn = attr_name
```

ipa component, BZ#983237

ipa-adtrust-install, an Identity Management Active Directory Trust configuration tool, does not explicitly specify authentication mechanism when performing Active Directory Trust configuration changes. When the user specifies the default LDAP authentication mechanism other than the expected default (for example, by setting the SASL_MECH configuration option to GSSAPI in the LDAP configuration file for the root user, **.ldaprc**), **ipa-adtrust-install** will not use the expected authentication mechanism and will fail to configure some of the parts of the Active Directory Integration feature, a crash of samba daemon (smbd) can occur or the user will be unable to use the feature. To work around this problem, remove any user default settings related to LDAP authentication mechanism from the **.ldaprc** file. The **ipa-adtrust-install** installer will then successfully configure the Active Directory integration feature.

ipa component, BZ#894388

The Identity Management installer configures all integrated services to listen on all interfaces. The administrator has no means to instruct the Identity Management installer to listen only on chosen interfaces even though the installer requires a valid interface IP address as one installation parameter. To work around this problem, change service configuration after Identity Management installation.

ipa component, BZ#894378

Identity Management LDAP permission manipulation plugin validates subtree and filter permission specifiers as mutually exclusive even though it is a valid combination in the underlying LDAP Access Control Instruction (ACI). Permissions with filter and subtree specifiers can be neither created nor

modified. This affects for example the **Add Automount Keys** permission which cannot be modified.

ipa component, BZ#817080

In some cases the certificates tracked by **certmonger** are not cleared when running the **ipa-server-install --uninstall** command. This will cause a subsequent re-installation to fail with an unexpected error.

sssd component, BZ#892604

The **ssh_cache** utility sets the DEBUG level after it processes the command-line parameters. If the command-line parameters cannot be processed, the utility prints DEBUG lines that are not supposed to be printed by default. To avoid this, correct parameters must be used.

sssd component, BZ#891647

It is possible to specify the **enumerate=true** value in the **sssd.conf** file to access all users in the system. However, using **enumerate=true** is not recommended in large environments as this can lead to high CPU consumption. As a result, operations like login or logout can be slowed down.

ipa component, BZ#888579

The Identity Management server processes Kerberos Password Expiration Time field as a 32-bit integer. If Maximum Lifetime of a user password in Identity Management Password Policy is set to a value causing the resulting Kerberos Password Expiration Time timestamp to exceed 32 bits and to overflow, the passwords that are being changed are configured with an expiration time that lies in the past and are always rejected. To ensure that new user passwords are valid and can be changed properly, do not set password Maximum Lifetime in Identity Management Password Policy to values that would cause the Kerberos Password Expiration Time timestamp to exceed 32 bits; that is, passwords that would expire after 2038-01-19. At the moment, recommended values for the Maximum Lifetime field are numbers lower than 9000 days.

sssd component, BZ#785877

When reconnecting to an LDAP server, SSSD does not check it was re-initialized during the downtime. If the server was re-initialized during the downtime and was filled with completely different data, SSSD does not update its database. As a consequence, the user can get invalid information from SSSD. To work around this problem:

1. stop SSSD before reconnecting to the re-initialized server;
2. clear the SSSD caches manually before reconnecting;
3. start SSSD.

krb5 component

In environments where entropy is scarce, the **kadmind** tool can take longer to initialize after startup than it did in previous releases as it attempts to read data from the **/dev/random** file and seed its internal random number generator (RNG). Clients which attempt to connect to the **kadmin** service can time out and fail with a GSS-API or Kerberos error. After the service completely finishes initializing itself, it will process messages received from now-disconnected clients and can log clock-skew or decrypt-integrity-check-failed errors for those connections. To work around this problem, use a service such as **rngd** to seed the system RNG using hardware sources of entropy.

ipa component, BZ#887193

The Identity Management server in Red Hat Enterprise Linux 6.3 introduced a technical preview of

SELinux user mapping feature, which enabled a mapping of SELinux users to users managed by the Identity Management based on custom rules. However, the default configured SELinux user (`guest_u:s0`) used when no custom rule matches is too constraining. An Identity Management user authenticating to Red Hat Enterprise Linux 6.5 can be assigned the too constraining SELinux user in which case a login through graphical session would always fail. To work around this problem, change a too constraining default SELinux user in the Identity Management server from `guest_u:s0` to a more relaxed value `unconfined_u:s0-s0:c0.c1023`:

```
kinit admin
ipa config-mod ipaselinusermapdefault=unconfined_u:s0-s0:c0.c1023
```

An unconfined SELinux user will be now assigned to the Identity Management user by default, which will allow the user to successfully authenticate through graphical interface.

ipa component, BZ#761574

When attempting to view a host in the web UI, the following message can appear:

```
Certificate operation cannot be completed: Unable to communicate with
CMS (Unauthorized)
```

Attempting to delete installed certificates through the web UI or command-line interface can fail with the same error message. To work around this problem, run the following command:

```
~]# yum downgrade ipa-server libipa_hbac libipa_hbac-python ipa-python
ipa-client ipa-admintools ipa-server-selinux
```

ipa component

When upgrading the ipa-server package using **anaconda**, the following error message is logged in the `upgrade.log` file:

```
/sbin/restorecon: lstat(/var/lib/pki-ca/publish*) failed: No such file
or directory
```

This problem does not occur when using **yum**.

sssd component

In the Identity Manager subdomain code, a User Principal Name (UPN) is by default built from the SAM Account Name and Active Directory trust users, that is `user@DOMAIN`. The UPN can be changed to differ from the UPN in Active Directory, however only the default format, `user@DOMAIN`, is supported.

sssd component, BZ#805921

Sometimes, group members may not be visible when running the `getent group groupname` command. This can be caused by an incorrect `ldap_schema` in the `[domain/DOMAINNAME]` section of the `sssd.conf` file. **SSSD** supports three LDAP schema types: RFC 2307, RFC 2307bis, and IPA. By default, **SSSD** uses the more common RFC 2307 schema. The difference between RFC 2307 and RFC 2307bis is the way which group membership is stored in the LDAP server. In an RFC 2307 server, group members are stored as the multi-valued `memberuid` attribute which contains the name of the users that are members. In an RFC2307bis server, group members are stored as the multi-valued attribute `member` (or sometimes `uniqueMember`) which contains the DN of the user or group that is a member of this group. RFC2307bis allows nested groups to be maintained as well.

When encountering this problem:

- add `ldap_schema = rfc2307bis` in the `sssd.conf` file,
- delete the `/var/lib/sss/db/cache_DOMAINNAME.ldb` file,
- and restart **SSSD**.

If the workaround does not work, add `ldap_group_member = uniqueMember` in the `sssd.conf` file, delete the cache file and restart SSSD.

Identity Management component, BZ#826973

When Identity Management is installed with its CA certificate signed by an external CA, the installation is processed in 2 stages. In the first stage, a CSR is generated to be signed by an external CA. The second stage of the installation then accepts a file with the new signed certificate for the Identity Management CA and a certificate of the external CA. During the second stage of the installation, a signed Identity Management CA certificate subject is validated. However, there is a bug in the certificate subject validation procedure and its default value (`O=$REALM`, where `$REALM` is the realm of the new Identity Management installation) is never pulled. Consequently, the second stage of the installation process always fails unless the `--subject` option is specified. To work around this issue, add the following option for the second stage of the installation: `--subject "O=$REALM"` where `$REALM` is the realm of the new Identity Management installation. If a custom subject was used for the first stage of the installation, use its value instead. Using this work around, the certificate subject validation procedure succeeds and the installation continues as expected.

Identity Management component, BZ#822350

When a user is migrated from a remote LDAP, the user's entry in the Directory Server does not contain Kerberos credentials needed for a Kerberos login. When the user visits the password migration page, Kerberos credentials are generated for the user and logging in via Kerberos authentication works as expected. However, Identity Management does not generate the credentials correctly when the migrated password does not follow the password policy set on the Identity Management server. Consequently, when the password migration is done and a user tries to log in via Kerberos authentication, the user is prompted to change the password as it does not follow the password policy, but the password change is never successful and the user is not able to use Kerberos authentication. To work around this issue, an administrator can reset the password of a migrated user with the `ipa passwd` command. When reset, user's Kerberos credentials in the Directory Server are properly generated and the user is able to log in using Kerberos authentication.

Identity Management component

In the Identity Management webUI, deleting a DNS record may, under some circumstances, leave it visible on the page showing DNS records. This is only a display issue and does not affect functionality of DNS records in any way.

Identity Management component, BZ#790513

The `ipa-client` package does not install the `policycoreutils` package as its dependency, which may cause install/uninstall issues when using the `ipa-client-install` setup script. To work around this issue, install the `policycoreutils` package manually:

```
~]# yum install policycoreutils
```

Identity Management component, BZ#813376

Updating the Identity Management LDAP configuration via the **ipa-ldap-updater** fails with a traceback error when executed by a non-root user due to the SASL EXTERNAL bind requiring root privileges. To work around this issue, run the aforementioned command as the root user.

Identity Management component, BZ#794882

With netgroups, when adding a host as a member that Identity Management does not have stored as a host already, that host is considered to be an external host. This host can be controlled with netgroups, but Identity Management has no knowledge of it. Currently, there is no way to use the **netgroup-find** option to search for external hosts.

Also, note that when a host is added to a netgroup as an external host, rather than being added in Identity Management as an external host, that host is not automatically converted within the netgroup rule.

Identity Management component, BZ#786629

Because a permission does not provide write access to an entry, delegation does not work as expected. The 389 Directory Server (**389-ds**) distinguishes access between entries and attributes. For example, an entry can be granted add or delete access, whereas an attribute can be granted read, search, and write access. To grant write access to an entry, the list of writable attributes needs to be provided. The **filter**, **subtree**, and other options are used to target those entries which are writable. Attributes define which part(s) of those entries are writable. As a result, the list of attributes will be writable to members of the permission.

sssd component, BZ#808063

The manpage entry for the **ldap_disable_paging** option in the **sssd-ldap** man page does not indicate that it accepts the boolean values True or False, and defaulting to False if it is not explicitly specified.

Identity Management component, BZ#812127

Identity Management relies on the LDAP schema to know what type of data to expect in a given attribute. If, in certain situations (such as replication), data that does not meet those expectations is inserted into an attribute, Identity Management will not be able to handle the entry, and LDAP tools have to be used to manually clean up that entry.

Identity Management component, BZ#812122

Identity Management **sudo** commands are not case sensitive. For example, executing the following commands will result in the latter one failing due to the case insensitivity:

```
~]$ ipa sudocmd-add /usr/bin/X
:
~]$ ipa sudocmd-add /usr/bin/x
ipa: ERROR: sudo command with name "/usr/bin/x" already exists
```

Identity Management component

When an Identity Management server is installed with a custom hostname that is not resolvable, the **ipa-server-install** command should add a record to the static hostname lookup table in **/etc/hosts** and enable further configuration of Identity Management integrated services. However, a record is not added to **/etc/hosts** when an IP address is passed as an CLI option and not interactively. Consequently, Identity Management installation fails because integrated services that are being configured expect the Identity Management server hostname to be resolvable. To work around this issue, complete one of the following:

- Run the `ipa-server-install` without the `--ip-address` option and pass the IP address interactively.
- Add a record to `/etc/hosts` before the installation is started. The record should contain the Identity Management server IP address and its full hostname (the `hosts(5)` man page specifies the record format).

As a result, the Identity Management server can be installed with a custom hostname that is not resolvable.

sssd component

Upgrading SSSD from the version provided in Red Hat Enterprise Linux 6.1 to the version shipped with Red Hat Enterprise Linux 6.2 may fail due to a bug in the dependent library `libldb`. This failure occurs when the SSSD cache contains internal entries whose distinguished name contains the `\,` character sequence. The most likely example of this is for an invalid `memberUID` entry to appear in an LDAP group of the form:

```
memberUID: user1,user2
```

`memberUID` is a multi-valued attribute and should not have multiple users in the same attribute.

If the upgrade issue occurs, identifiable by the following debug log message:

```
(Wed Nov  2 15:18:21 2011) [sssd] [ldb] (0): A transaction is still
active in
ldb context [0xaa0460] on /var/lib/sss/db/cache_<DOMAIN>.ldb
```

remove the `/var/lib/sss/db/cache_<DOMAIN>.ldb` file and restart SSSD.



WARNING

Removing the `/var/lib/sss/db/cache_<DOMAIN>.ldb` file purges the cache of all entries (including cached credentials).

sssd component, [BZ#751314](#)

When a group contains certain incorrect multi-valued `memberUID` values, SSSD fails to sanitize the values properly. The `memberUID` value should only contain one username. As a result, SSSD creates incorrect users, using the broken `memberUID` values as their usernames. This, for example, causes problems during cache indexing.

Identity Management component

Two Identity Management servers, both with a CA (Certificate Authority) installed, use two replication agreements. One is for user, group, host, and other related data. Another replication agreement is established between the CA instances installed on the servers. If the CA replication agreement is broken, the Identity Management data is still shared between the two servers, however, because there is no replication agreement between the two CAs, issuing a certificate on one server will cause the other server to not recognize that certificate, and vice versa.

Identity Management component

The Identity Management (ipa) package cannot be build with a **6ComputeNode** subscription.

sssd component, BZ#741264

Active Directory performs certain LDAP referral-chasing that is incompatible with the referral mechanism included in the **openldap** libraries. Notably, Active Directory sometimes attempts to return a referral on an LDAP bind attempt, which used to cause a hang, and is now denied by the **openldap** libraries. As a result, SSSD may suffer from performance issues and occasional failures resulting in missing information.

To work around this issue, disable referral-chasing by setting the following parameter in the `[domain/DOMAINNAME]` section of the `/etc/sss/sss.conf` file:

```
ldap_referrals = false
```

6.10. DEVICES

kernel component

When using large block size (1MB), the tape driver sometimes returns an EBUSY error. To work around this problem, use a smaller block size, that is 256KB.

kernel component

On some of the older Broadcom tg3 devices, the default Maximum Read Request Size (MRRS) value of 512 byte is known to cause lower performance. It is because these devices perform direct memory access (DMA) requests serially. 1500-byte ethernet packet will be broken into 3 PCIE read requests using 512 byte MRRS. When using a higher MRRS value, the DMA transfer can be faster as fewer requests will be needed. However, the MRRS value is meant to be tuned by system software and not by the driver. PCIE Base spec 3.0 section 7.8.4 contains an implementation note that illustrates how system software might tune the MRRS for all devices in the system. As a result, Broadcom modified the tg3 driver to remove the code that sets the MRRS to 4K bytes so that any value selected by system software (BIOS) will be preserved.

kernel component

The Brocade BFA Fibre Channel and FCoE driver does not currently support dynamic recognition of Logical Unit addition or removal using the **sg3_utils** utilities (for example, the **sg_scan** command) or similar functionality. Please consult Brocade directly for a Brocade equivalent of this functionality.

kexec-tools component

Starting with Red Hat Enterprise Linux 6.0 and later, kexec kdump supports dumping core to the Btrfs file system. However, note that because the **findfs** utility in **busybox** does not support Btrfs yet, **UUID/LABEL** resolving is not functional. Avoid using the **UUID/LABEL** syntax when dumping core to Btrfs file systems.

trace-cmd component

The **trace-cmd** service does not start on 64-bit PowerPC and IBM System z systems because the **sys_enter** and **sys_exit** events do not get enabled on the aforementioned systems.

trace-cmd component

trace-cmd's subcommand, **report**, does not work on IBM System z systems. This is due to the fact that the **CONFIG_FTRACE_SYSCALLS** parameter is not set on IBM System z systems.

libfprint component

Red Hat Enterprise Linux 6 only has support for the first revision of the UPEK Touchstrip fingerprint reader (USB ID 147e:2016). Attempting to use a second revision device may cause the fingerprint reader daemon to crash. The following command returns the version of the device being used in an individual machine:

```
~]$ lsusb -v -d 147e:2016 | grep bcdDevice
```

kernel component

The Emulex Fibre Channel/Fibre Channel-over-Ethernet (FCoE) driver in Red Hat Enterprise Linux 6 does not support DH-CHAP authentication. DH-CHAP authentication provides secure access between hosts and mass storage in Fibre-Channel and FCoE SANs in compliance with the FC-SP specification. Note, however that the Emulex driver (**lpfc**) does support DH-CHAP authentication on Red Hat Enterprise Linux 5, from version 5.4. Future Red Hat Enterprise Linux 6 releases may include DH-CHAP authentication.

kernel component

The recommended minimum HBA firmware revision for use with the **mpt2sas** driver is "Phase 5 firmware" (that is, with version number in the form **05.xx.xx.xx**). Note that following this recommendation is especially important on complex SAS configurations involving multiple SAS expanders.

6.11. KERNEL

kernel component

Sun Fire X4500 data server enumerates the e1000 card with Peripheral Component Interconnect Extended (PCI-X) and enables 64-bit direct memory access (DMA), however, 64-bit DMA is not fully supported on this hardware. If possible, disable 64-bit DMA in BIOS.

grubby component

Use of multiboot images makes discerning different image types problematic during kernel updates. As a consequence, using the **tboot** package and multiple types of kernels at the same time does not work properly. If, for example, **tboot** is in use and the **kernel-debug** package is installed, bootloader configuration can sometimes reflect an incorrect image list. To avoid this, do not use the **kernel-debug** on a system utilizing **tboot**, or vice versa. If such a situation is unavoidable, manually verify that the bootloader configuration is reasonable after each update before rebooting.

kexec-tools component

When the debug kernel is installed and also used as the Red Hat Enterprise Linux **kdump** kernel, the reserved **kdump** memory must be increased to a minimum of 256 MB. To assure this setting, start the **system-config-kdump** tool, modify the **kdump** memory, and reboot your Linux instance. Alternatively, you can configure a particular kernel that is always used as the **kdump** kernel, independently of the running kernel. For more information, consult the [Red Hat Enterprise Linux 6 Deployment Guide](#).

kernel component

Red Hat Enterprise Linux 6.4 changed the maximum read/write socket memory default value to be higher, allowing for better performance on some machines. It was observed that if the values of `mem_max` are not symmetrical between two machines, the performance can be negatively affected. To work around this problem, adjust the value of `mem_max` to be equal across all Red Hat Enterprise Linux systems in the network.

kabi-whitelists component

The `vxfs` module might not work properly on Red Hat Enterprise Linux 6.4 and later because of the broken `radix_tree_gang_lookup_slot` symbol. Consult Symantec should you require a workaround for this issue.

kernel component

Enabling TCP Segmentation Offload (TSO) on TAP interface may cause low throughput when the uplink is a high-speed interface. To improve throughput, turn off TSO on the tap interface of the virtual machine.

kernel component

When using Chelsio's iSCSI HBAs for an iSCSI root partition, the first boot after install fails. This occurs because Chelsio's iSCSI HBA is not properly detected. To work around this issue, users must add the `iscsi_firmware` parameter to grub's kernel command line. This will signal to dracut to boot from the iSCSI HBA.

kernel component

The installation of Red Hat Enterprise Linux 6.3 i386 and later may occasionally fail. To work around this issue, add the following parameter to the kernel command line:

```
vmalloc=256MB
```

kernel component

If a device reports an error, while it is opened (via the `open(2)` system call), then the device is closed (via the `close(2)` system call), and the `/dev/disk/by-id` link for the device may be removed. When the problem on the device that caused the error is resolved, the `by-id` link is not re-created. To work around this issue, run the following command:

```
~]# echo 'change' > /sys/class/block/sdX/uevent
```

kernel component

When an HBA that uses the `mpt2sas` driver is connected to a storage using an SAS switch LSI SAS 6160, the driver may become unresponsive during Controller Fail Drive Fail (CFDF) testing. This is due to faulty firmware that is present on the switch. To fix this issue, use a newer version (14.00.00.00 or later) of firmware for the LSI SAS 6160 switch.

kernel component, BZ#745713

In some cases, Red Hat Enterprise Linux 6 guests running fully-virtualized under Red Hat Enterprise Linux 5 experience a time drift or fail to boot. In other cases, drifting may start after migration of the virtual machine to a host with different speed. This is due to limitations in the Red Hat Enterprise Linux 5 Xen hypervisor. To work around this, add the `nohpet` parameter or, alternatively, the

`clocksource=jiffies` parameter to the kernel command line of the guest. Or, if running under Red Hat Enterprise Linux 5.7 or newer, locate the guest configuration file for the guest and add the `hpet=0` parameter in it.

kernel component

On some systems, Xen full-virt guests may print the following message when booting:

```
WARNING: BIOS bug: CPU MTRRs don't cover all of memory, losing
<number>MB of RAM
```

It is possible to avoid the memory trimming by using the `disable_mtrr_trim` kernel command line option.

kernel component

The `perf record` command becomes unresponsive when specifying a tracepoint event and a hardware event at the same time.

kernel component

On 64-bit PowerPC, the following command may cause kernel panic:

```
~]# ./perf record -agT -e sched:sched_switch -F 100 -- sleep 3
```

kernel component

Applications are increasingly using more than 1024 file descriptors. It is not recommended to increase the default soft limit of file descriptors because it may break applications that use the `select()` call. However, it is safe to increase the default hard limit; that way, applications requiring a large amount of file descriptors can increase their soft limit without needing root privileges and without any user intervention.

kernel component

In network only use of Brocade Converged Network Adapters (CNAs), switches that are not properly configured to work with Brocade FCoE functionality can cause a continuous linkup/linkdown condition. This causes continuous messages on the host console:

```
bfa xxxx:xx:xx.x: Base port (WWN = xx:xx:xx:xx:xx:xx:xx:xx) lost fabric
connectivity
```

To work around this issue, unload the Brocade `bfa` driver.

kernel component

In Red Hat Enterprise Linux 6, a legacy bug in the PowerEdge Expandable RAID Controller 5 (PERC5) which causes the `kdump` kernel to fail to scan for `scsi` devices. It is usually triggered when a large amounts of I/O operations are pending on the controller in the first kernel before performing a `kdump`.

kernel component, [BZ#679262](#)

In Red Hat Enterprise Linux 6.2 and later, due to security concerns, addresses in `/proc/kallsyms` and `/proc/modules` show all zeros when accessed by a non-root user.

kernel component

Superfluous information is displayed on the console due to a correctable machine check error occurring. This information can be safely ignored by the user. Machine check error reporting can be disabled by using the **nomce** kernel boot option, which disables machine check error reporting, or the **mce=ignore_ce** kernel boot option, which disables correctable machine check error reporting.

kernel component

The order in which PCI devices are scanned may change from one major Red Hat Enterprise Linux release to another. This may result in device names changing, for example, when upgrading from Red Hat Enterprise Linux 5 to 6. You must confirm that a device you refer to during installation, is the intended device.

One way to assure the correctness of device names is to, in some configurations, determine the mapping from the controller name to the controller's PCI address in the older release, and then compare this to the mapping in the newer release, to ensure that the device name is as expected.

The following is an example from `/var/log/messages`:

```
kernel: cciss0: <0x3230> at PCI 0000:1f:00.0 IRQ 71 using DAC
...
kernel: cciss1: <0x3230> at PCI 0000:02:00.0 IRQ 75 using DAC
```

If the device name is incorrect, add the **pci=bfsort** parameter to the kernel command line, and check again.

kernel component

The minimum firmware version for NIC adapters managed by **netxen_nic** is 4.0.550. This includes the boot firmware which is flashed in option ROM on the adapter itself.

kernel component

High stress on 64-bit IBM POWER series machines prevents `kdump` from successfully capturing the **vmcore**. As a result, the second kernel is not loaded, and the system becomes unresponsive.

kernel component

Triggering `kdump` to capture a **vmcore** through the network using the Intel 82575EB ethernet device in a 32 bit environment causes the networking driver to not function properly in the `kdump` kernel, and prevent the **vmcore** from being captured.

kernel component

Memory Type Range Register (MTRR) setup on some hyperthreaded machines may be incorrect following a suspend/resume cycle. This can cause graphics performance (specifically, scrolling) to slow considerably after a suspend/resume cycle.

To work around this issue, disable and then re-enable the hyperthreaded sibling CPUs around suspend/resume, for example:

```
#!/bin/sh
# Disable hyper-threading processor cores on suspend and hibernate, re-
enable
# on resume.
# This file goes into /etc/pm/sleep.d/
```

```
case $1 in
    hibernate|suspend)
        echo 0 > /sys/devices/system/cpu/cpu1/online
        echo 0 > /sys/devices/system/cpu/cpu3/online
        ;;
    thaw|resume)
        echo 1 > /sys/devices/system/cpu/cpu1/online
        echo 1 > /sys/devices/system/cpu/cpu3/online
        ;;
esac
```

kernel component

In Red Hat Enterprise Linux 6.2, **nmi_watchdog** registers with the **perf** subsystem. Consequently, during boot, the **perf** subsystem grabs control of the performance counter registers, blocking OProfile from working. To resolve this, either boot with the **nmi_watchdog=0** kernel parameter set, or run the following command to disable it at run time:

```
echo 0 > /proc/sys/kernel/nmi_watchdog
```

To re-enable **nmi-watchdog**, use the following command

```
echo 1 > /proc/sys/kernel/nmi_watchdog
```

kernel component, [BZ#603911](#)

Due to the way **ftrace** works when modifying the code during start-up, the NMI watchdog causes too much noise and **ftrace** can not find a quiet period to instrument the code. Consequently, machines with more than 512 CPUs will encounter issues with the NMI watchdog. Such issues will return error messages similar to **BUG: NMI Watchdog detected LOCKUP** and have either **ftrace_modify_code** or **ipi_handler** in the backtrace. To work around this issue, disable NMI watchdog by setting the **nmi_watchdog=0** kernel parameter, or using the following command at run time:

```
echo 0 > /proc/sys/kernel/nmi_watchdog
```

kernel component

On 64-bit POWER systems the EHEA NIC driver will fail when attempting to dump a **vmcore** via NFS. To work around this issue, utilize other kdump facilities, for example dumping to the local file system, or dumping over SSH.

kernel component, [BZ#587909](#)

A BIOS emulated floppy disk might cause the installation or kernel boot process to hang. To avoid this, disable emulated floppy disk support in the BIOS.

kernel component

The preferred method to enable **nmi_watchdog** on 32-bit x86 systems is to use either **nmi_watchdog=2** or **nmi_watchdog=lpic** parameters. The parameter **nmi_watchdog=1** is not supported.

kernel component

The kernel parameter, **pci=noioapicquirk**, is required when installing the 32-bit variant of Red Hat Enterprise Linux 6 on HP xw9300 workstations. Note that the parameter change is not required when installing the 64-bit variant.

6.12. DESKTOP

gnome-panel component, BZ#1017631

The **gnome-panel** utility can sometimes terminate unexpectedly on 64-bit PowerPC architecture using the XDMCP protocol.

xorg-x11-drv-intel component, BZ#889574

Red Hat Enterprise Linux 6 graphics stacks does not support NVIDIA Optimus hardware configurations. On laptops with both Intel and NVIDIA GPUs, some or all external video ports may not function correctly when using the Intel GPU. If external video ports are needed, configure the BIOS to use the NVIDIA GPU instead of the Intel GPU if possible.

xorg-x11-drv-synaptics component, BZ#873721

Two-finger scrolling is default for devices that announce two-finger capability. However, on certain machines, although the touchpad announces two-finger capability, events generated by the device only contain a single finger position at a time and two-finger scrolling therefore does not work. To work around this problem, use edge scrolling instead.

firefox component

In certain environments, storing personal Firefox configuration files (~/.mozilla/) on an NFS share, such as when your home directory is on a NFS share, led to Firefox functioning incorrectly, for example, navigation buttons not working as expected, and bookmarks not saving. This update adds a new configuration option, `storage.nfs_filesystem`, that can be used to resolve this issue. If you experience this issue:

1. Start **Firefox**.
2. Type **about:config** into the URL bar and press the **Enter** key.
3. If prompted with "This might void your warranty!", click the **I'll be careful, I promise!** button.
4. Right-click in the **Preference Name** list. In the menu that opens, select **New** → **Boolean**.
5. Type "storage.nfs_filesystem" (without quotes) for the preference name and then click the **OK** button.
6. Select **true** for the boolean value and then press the **OK** button.

wacomcpl component, BZ#769466

The `wacomcpl` package has been deprecated and has been removed from the package set. The `wacomcpl` package provided graphical configuration of Wacom tablet settings. This functionality is now integrated into the GNOME Control Center.

acroread component

Running a AMD64 system without the `sssd-client.i686` package installed, which uses SSSD for getting information about users, causes **acoread** to fail to start. To work around this issue, manually install the `sssd-client.i686` package.

kernel component, BZ#681257

With newer kernels, such as the kernel shipped in Red Hat Enterprise Linux 6.1, Nouveau has corrected the Transition Minimized Differential Signaling (TMDS) bandwidth limits for pre-G80 NVIDIA chipsets. Consequently, the resolution auto-detected by X for some monitors may differ from that used in Red Hat Enterprise Linux 6.0.

fprintd component

When enabled, fingerprint authentication is the default authentication method to unlock a workstation, even if the fingerprint reader device is not accessible. However, after a 30 second wait, password authentication will become available.

evolution component

Evolution's IMAP backend only refreshes folder contents under the following circumstances: when the user switches into or out of a folder, when the auto-refresh period expires, or when the user manually refreshes a folder (that is, using the menu item **Folder** → **Refresh**). Consequently, when replying to a message in the Sent folder, the new message does not immediately appear in the Sent folder. To see the message, force a refresh using one of the methods describe above.

anaconda component

The clock applet in the GNOME panel has a default location of Boston, USA. Additional locations are added via the applet's preferences dialog. Additionally, to change the default location, left-click the applet, hover over the desired location in the **Locations** section, and click the **Set . . .** button that appears.

xorg-x11-server component, BZ#623169

In some multi-monitor configurations (for example, dual monitors with both rotated), the cursor confinement code produces incorrect results. For example, the cursor may be permitted to disappear off the screen when it should not, or be prevented from entering some areas where it should be allowed to go. Currently, the only workaround for this issue is to disable monitor rotation.

6.13. TOOLS

ssh-keygen component

The following example in the description of the `-V` option in the `ssh-keygen(1)` manual page is incorrect:

```
“-4w:+4w” (valid from four weeks ago to four weeks from now)
```

If you set a date range in this format, the certificate is valid from four weeks ago until now.

perl-www-curl component

Attempting to access the `CURLINFO_PRIVATE` value can cause **curl** to terminate unexpectedly with a segmentation fault.

freerpd component, BZ#988277

The ALSA plug-in is not supported in Red Hat Enterprise Linux 6. Instead of the ALSA plug-in, use the pulseaudio plug-in. To enable it, use the `--plugin rplsnd` option with the `xfreerdp` command without specifying which plug-in should be used; the pulseaudio plug-in will be used automatically in this case.

coolkey component, BZ#906537

Personal Identity Verification (PIV) Endpoint Cards which support both CAC and PIV interfaces might not work with the latest **coolkey** update; some signature operations like PKINIT can fail. To work around this problem, downgrade **coolkey** to the version shipped with Red Hat Enterprise Linux 6.3.

libreport component

Even if the stored credentials are used, the **report-gtk** utility can report the following error message:

```
Wrong settings detected for Red Hat Customer Support [..]
```

To work around this problem, close the dialog window; the **Login=<rhel-user>** and **Password=<rhel-password>** credentials in the `/etc/libreport/plugins/rhtsupport.conf` will be used in the same way they are used by **report-rhtsupport**.

For more information, refer to [this](#) Knowledge Base article.

vlock component

When a user password is used to lock a console with **vlock**, the console can only be unlocked with the user password, not the root password. That is, even if the first inserted password is incorrect, and the user is prompted to provide the root password, entering the root password fails with an error message.

libreoffice component

Libreoffice contains a number of harmless files used for testing purposes. However, on Microsoft Windows system, these files can trigger false positive alerts on various anti-virus software, such as Microsoft Security Essentials. For example, the alerts can be triggered when scanning the Red Hat Enterprise Linux 6 ISO file.

gnome-power-manager component

When the computer runs on battery, custom brightness level is not remembered and restored if power saving features like "dim display when idle" or "reduce backlight brightness when idle" are enabled.

rsyslog component

rsyslog does not reload its configuration after a **SIGHUP** signal is issued. To reload the configuration, the **rsyslog** daemon needs to be restarted:

```
~]# service rsyslog restart
```

6.14. DOCUMENTATION

release-notes component

The Release Notes document included in Red Hat Enterprise Linux 6.5 and available on the [Customer Portal](#) contains incorrectly lists information about the FSTEK certification in all languages. Please consult the online English version of the [Release Notes](#), which is the latest and most up-to-

date version.

release-notes component

The Benagali (bn-IN) and Simplified Chinese (zh-CN) translations of the Release Notes included in Red Hat Enterprise Linux 6.5 and on the [Customer Portal](#) contain several untranslated strings.

CHAPTER 7. NEW PACKAGES

7.1. RHEA-2013:1625 — NEW PACKAGES: FREERDP

New freerdp packages are now available for Red Hat Enterprise Linux 6.

FreeRDP is a free implementation of the Remote Desktop Protocol (RDP), released under the Apache license. The xfreerdp client can connect to RDP servers such as Microsoft Windows machines, xrdp and VirtualBox.

This enhancement update adds the freerdp packages to Red Hat Enterprise Linux 6. (BZ#951696)

All users who require freerdp are advised to install these new packages.

7.2. RHBA-2013:1607 — NEW PACKAGES: GCC-LIBRARIES

New gcc-libraries packages are now available for Red Hat Enterprise Linux 6.

The new gcc-libraries packages contain various GCC runtime libraries, such as libatomic and libitm. In Red Hat Enterprise Linux 5.9, libitm was a separate package that included the libitm library. The libitm package is now deprecated and replaced by the gcc-libraries packages.

This enhancement update adds the gcc-libraries packages to Red Hat Enterprise Linux 6. (BZ#906241)

All users who require gcc-libraries are advised to install these new packages.

7.3. RHEA-2013:1728 — NEW PACKAGES: OPENHPI32

New openhpi32 packages are now available for Red Hat Enterprise Linux 6.

OpenHPI provides an open source implementation of the Service Availability Forum (SAF) Hardware Platform Interface (HPI). HPI is an abstracted interface for managing computer hardware, typically chassis- and rack-based servers. HPI includes resource modeling; access to and control over sensor, control, watchdog, and inventory data associated with resources; abstracted System Event Log interfaces; hardware events and alarms; and a managed hot swap interface. This is version 3.2 of the OpenHPI project.

This enhancement update adds the openhpi32 packages to Red Hat Enterprise Linux 6. (BZ#927897)

All users who require openhpi32 are advised to install these new packages.

7.4. RHEA-2013:1626 — NEW PACKAGES: P11-KIT

New p11-kit packages are now available for Red Hat Enterprise Linux 6.

The p11-kit package provides a mechanism to manage PKCS#11 modules. The p11-kit-trust subpackage includes a PKCS#11 trust module that provides certificate anchors and black lists based on configuration files.

This enhancement update adds the p11-kit packages to Red Hat Enterprise Linux 6. (BZ#915798)

* Red Hat Enterprise Linux 6.5 provides the p11-kit package to implement the Shared System Certificates feature. If enabled by the administrator, it ensures system-wide trust store of static data that is used by crypto toolkits as input for certificate trust decisions. (BZ#977886)

These new packages had several bugs fixed during testing:

- * Support for using the freebl3 library for the SHA1 and MD5 cryptographic hash functions has been added even though the hashing is done in a strictly non-cryptographic context. (BZ#983384)
- * All file handles opened by p11-kit are created with the O_CLOEXEC flag, so that they are automatically closed on the execve() function and do not leak to subprocesses. (BZ#984986)
- * When expanding the "\$HOME" variable or the "~/ " path for SUID and SGID programs, the expand_home() function returns NULL. This change allows for avoiding vulnerabilities that could occur if SUID or SGID programs accidentally trusted this environment. Also, documentation concerning the fact that user directories are not read for SUID/SGID programs has been added. (BZ#985014)
- * Users need to use the standard environment \$TMPDIR variable for locating the temp directory. (BZ#985017)
- * If a critical module fails to initialize, module initialization stops and the user is informed about the failure. (BZ#985023)
- * The p11_kit_space_strlen() function returns a "0" value for empty strings. (BZ#985416)
- * Arguments of the size_t variable are correctly passed to the "p11_hash_xxx" functions. (BZ#985421)
- * Changes in the code ensures that the memdup() function is not called with a zero length or NULL pointers. (BZ#985433)

All users who require the Shared System Certificates feature are advised to install these new packages.

7.5. RHEA-2013:1621 — NEW PACKAGE: PS_MEM

A new ps_mem package is now available for Red Hat Enterprise Linux 6.

The ps_mem package provides a memory usage script written in Python that calculates how much RAM is used per program. The script automatically selects the most accurate method, which is available for a particular running kernel.

This enhancement update adds the ps_mem package to Red Hat Enterprise Linux 6. (BZ#962850)

All users who require ps_mem are advised to install this new package.

7.6. RHEA-2013:1642 — NEW PACKAGES: REDHAT-SUPPORT-LIB-PYTHON AND REDHAT-SUPPORT-TOOL

New redhat-support-lib-python and redhat-support-tool packages are now available for Red Hat Enterprise Linux 6.

The redhat-support-lib-python package provides a Python library that developers can use to easily write software solutions that leverage Red Hat Access subscription services.

The redhat-support-tool utility facilitates console-based access to Red Hat's subscriber services and gives Red Hat subscribers more venues for accessing the content and services available to them as Red Hat customers. Further, it enables our customers to integrate and automate their helpdesk services with our subscription services. The capabilities of this package include:

- * Red Hat Access Knowledge Base article and solution viewing from the console (formatted as man pages).
- * Viewing, creating, modifying, and commenting on customer support cases from the console. *

Attachment uploading directly to a customer support case or to <ftp://dropbox.redhat.com/> from the console. * Full proxy support (that is, FTP and HTTP proxies). * Easy listing and downloading of attachments in customer support cases from the console. * Red Hat Access Knowledge Base searching on query terms, log messages, and other parameters, and viewing search results in a selectable list. * Easy uploading of log files, text files, and other sources to the Red Hat Access automatic problem determination engine for diagnosis. * Various other support-related commands.

Detailed usage information for the tool can be found in the Red Hat Customer Portal at <https://access.redhat.com/site/articles/445443>

This enhancement update adds the `redhat-support-lib-python` and `redhat-support-tool` packages to Red Hat Enterprise Linux 6. (BZ#[987159](#), BZ#[869395](#), BZ#[880776](#), BZ#[987171](#), BZ#[987169](#), BZ#[987163](#))

All users who require `redhat-support-lib-python` and `redhat-support-tool` are advised to install these new packages.

7.7. RHEA-2013:1686 — NEW PACKAGE: SAPCONF

A new `sapconf` package is now available for Red Hat Enterprise Linux 6.

The `sapconf` package contains a script that checks the basic installation of Red Hat Enterprise Linux and modifies it according to SAP requirements. The script ensures that all necessary packages are installed and that configuration parameters are set correctly to run SAP software.

This enhancement update adds the `sapconf` package to Red Hat Enterprise Linux 6. This package is available through the "Red Hat Enterprise Linux for SAP Business Applications" channel. (BZ#[910838](#))

All users who running SAP software on Red Hat Enterprise Linux 6 are advised to install this new package.

7.8. RHEA-2013:1731 — NEW PACKAGES: SNAPPY

New `snappy` packages are now available for Red Hat Enterprise Linux 6.

`Snappy` is a compression and decompression library that aims for very high speeds and reasonable compression.

This enhancement update adds the `snappy` packages to Red Hat Enterprise Linux 6. (BZ#[903090](#))

All users who require `snappy` are advised to install these new packages.

7.9. RHEA-2013:1622 — NEW PACKAGES: XORG-X11-GLAMOR

New `xorg-x11-glamor` packages are now available for Red Hat Enterprise Linux 6.

The `glamor` module is an open-source 2D graphics common driver for the X Window System as implemented by X.org. It supports a variety of graphics chip sets which have OpenGL, EGL or GBM support.

This enhancement update adds the `xorg-x11-glamor` packages to Red Hat Enterprise Linux 6. The `glamor` library is provided to support new AMD GPU hardware and can be used by the DDX driver to implement acceleration using the OpenGL driver. Some new hardware, such as AMD HD7xxx Series, needs `glamor` for acceleration. (BZ#[962832](#))

All users who require `xorg-x11-glamor` are advised to install these new packages.

CHAPTER 8. UPDATED PACKAGES

8.1. ABRT

8.1.1. [RHBA-2013:1586](#) — [abrt, libreport and btparser bug fix and enhancement update](#)

Updated abrt, libreport, and btparser packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

ABRT is a tool to help users to detect defects in applications and to create a problem report with all the information needed by a maintainer to fix it. ABRT uses a plug-in system to extend its functionality.

The libreport libraries provide an API for reporting different problems in applications to different bug targets like Bugzilla, ftp, and trac.

The btparser utility is a backtrace parser and analyzer library, which works with backtraces produced by the GNU Project Debugger. It can parse a text file with a backtrace to a tree of C structures, allowing to analyze the threads and frames of the backtrace and process them.

Bug Fixes

[BZ#854668](#)

If the `/etc/abrt/abrt.conf` file was modified so that the "DumpLocation" and "WatchCrashdumpArchiveDir" variables referred to the same directory, the ABRT utility tried to process the files in that directory as both archives and new problem directories, which led to unpredictable results. With this update, ABRT refuses to start if such misconfiguration is detected.

[BZ#896090](#)

While creating a case, the `reporter-rhtsupport` utility sent the operation system (OS) version value which RHT customer center server did not accept. Consequently, a new case failed to be created and an error message was returned. With this update, suffixes such as "Beta" in the OS version value are not stripped, RHT customer center server accepts the version value, and a case is created.

[BZ#952773](#)

Prior to this update, the `abrt-watch-log` and `abrt-dump-oops` utilities were creating too many new problem directories when a kernel error occurred periodically. As a consequence, the user was flooded with problem reports and the `/var` partition could overflow. To fix this bug, `abrt-dump-oops` has been changed to ignore all additional problems for a few minutes after it sees 5 or more of them. As a result, the user is not flooded with problem reports.

Enhancements

[BZ#952704](#)

The Red Hat Support tool required an API for querying crashes caught by ABRT. With this update, python API for ABRT has been provided and it is now possible to use python API to query bugs caught by ABRT.

[BZ#961231](#)

There is a high probability that users who do not use the graphical environment (headless systems) will miss the problems detected by the ABRT utility. When the user installs the `abrt-console-`

notification packages, they now see a warning message in the console regarding new problems detected since the last login.

All users of `abrt`, `libreport` and `btparser` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.2. ANACONDA

8.2.1. RHBA-2013:1588 — anaconda bug fix and enhancement update

Updated `anaconda` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Bug Fixes

BZ#818233

Previously, **anaconda** did not recognize a DVD ISO image written to a USB drive as a source repository for installation because this device does not have partitions. Consequently, the ISO acted like a `boot.iso` and it was not possible to install packages included in it. With this update, **anaconda** has been modified to include devices with ISO 9660 formatting, and to configure any device as a source repository if this device contains the `/repodata/repomd.xml` file. As a result, **anaconda** now recognizes ISO on USB as expected.

BZ#845572

Prior to this update, the **anaconda loader** command created the `/etc/sysconfig/network` file by renaming a new temporary file, which did not trigger the **NetworkManager**'s inotify mechanism. Consequently, a hostname set by the `network --hostname` kickstart option could be overridden by **NetworkManager** with hostname obtained through DHCP or DNS. With this update, **loader** has been modified to write new values directly into `/etc/sysconfig/network`. As a result, **NetworkManager** now accepts the hostname value specified in this file.

BZ#846336

Previously, **anaconda** did not attempt to use another loop device if the firstly selected one was already in use. Consequently, HDD ISO installation failed if loop devices were used in the kickstart `%pre` section. With this update, **anaconda** has been modified to use another loop device if the first one is already in use. As a result, HDD ISO installation works as expected.

BZ#847600

With this update, the `list-harddrives` command has been modified not to list the `/dev/srX` devices in its output.

BZ#851284

With this update, several typographical errors have been corrected in the *About LVM* dialog.

BZ#852523

When a different set of disks was used for the `clearpart --drives` and `part --ondisk` commands, a backtrace was returned. Consequently, installation did not finish successfully. With this update, only one set of disks is used with these commands. User must specify multiple disks with a single `clearpart` command, otherwise only the last `clearpart --drives` arguments is used.

BZ#859420

Prior to this update, when partitioning was incorrectly specified, the **No free space error** message was incorrectly shown instead of the appropriate **No free slots** dialog. With this update, the correct error message is displayed in case of incorrectly specified partitioning.

BZ#859569

Previously, **anaconda** in rescue mode unmounted the source ISO before searching for the **.discinfo** file. Consequently, the **stage2** parameter was loaded twice, increasing the boot time. With this update, **anaconda** has been modified to skip the check for **.discinfo** in rescue mode. As a result, **stage2** is only loaded once, as expected.

BZ#873281

Previously, when re-installing the system with already configured LVM raid1 volumes, **anaconda** terminated unexpectedly. This bug has been fixed, and **anaconda** no longer crashes in the aforementioned scenario.

BZ#875644

When a kickstart upgrade was performed on IBM System z architectures, **anaconda** shut down the system instead of rebooting even though the **reboot** command was present in the kickstart configuration. Consequently, a manual reboot was required. This update adds support for kickstart upgrades on System z, thus fixing this bug.

BZ#877852

Previously, when installing Red Hat Enterprise Linux 6 on a system with multiple disks and one or more of these disks contained the *PPC PReP Boot* partition, **anaconda** created an empty *PPC PReP Boot* partition on the selected installation disk, but stored necessary boot files in the already existing *PPC PReP Boot*. Consequently, the system failed to boot after the installation. With this update, **anaconda** has been modified to use the correct *PPC PReP Boot* for boot files, thus fixing this bug.

BZ#878907

The algorithm for calculating the swap size did not take into account the amount of space used for the installation. Consequently, even on small disks the installer created big swap space often leaving only insufficient amount of space for the rest of the system. This algorithm has been modified to register the amount of space used for the installation. As a result, smaller (10% of used disks' space) swap is created on machines with small disks leaving more space for the rest of the system.

BZ#880577

Previously, **anaconda** did not create partitions larger than 16TB on **XFS** filesystems. This bug has been fixed, and the official limit of 100TB is now used as accepted.

BZ#881005

Prior to this update, the **autopart** command did not function correctly with already defined prepboot partitions. Consequently, when using a **kickstart** file that contained the **part** command defining a prepboot partition followed by **autopart**, **anaconda** terminated unexpectedly with a segmentation fault. With this update, **autopart** has been modified to work correctly in the aforementioned configuration. As a result, the installation continues as expected.

BZ#882452

Previously, when configuring network devices within the **anaconda** GUI, devices using the FCoE network technology were automatically set not to be controlled by the **NetworkManager**.

Consequently, **NetworkManager** disabled these devices, causing previously connected FCoE SAN disks to disappear from the GUI. This bug has been fixed, and editing network device configuration in the GUI no longer disconnects previously set FCoE devices.

BZ#886020

Previously, **anaconda** did not return a warning message when using a raw partition for the / mount point without creating a new file system. With this update, **anaconda** has been modified to display a warning message in such scenario.

BZ#888292

Under certain circumstances, when managing partitions with the **anaconda** GUI, an unexpected loss of window focus occurred. With this update, the parent window setting has been modified, thus fixing this bug.

BZ#893849

With this update, several typographic and translation errors have been corrected in the Japanese locale in **anaconda**.

BZ#894050

Previously, **anaconda** created the `/etc/zipl.conf` configuration file using a set of default kernel parameters regardless of whether a fresh install or upgrade was performed. Consequently, kernel parameters added to `/etc/zipl.conf` by users were lost when upgrading IBM System z systems with **anaconda**. This update adds support for boot loader upgrades for systems with System z architecture. As a result, kernel parameters added by users to `/etc/zipl.conf` are preserved in the aforementioned scenario.

BZ#895098

Prior to this update, when attempting to install conflicting packages with the **anaconda** GUI, a misleading warning message was displayed. With this update, this message has been modified to inform about the package conflict.

BZ#895982

Physical-extents size less than 32MB on top of an MD physical volume leads **anaconda** to problems with calculating the capacity of a volume group. To work around this problem, use a physical-extent size of at least 32MB or leave free space (with size equal to doubled size of the physical-extent) when allocating logical volumes.

BZ#901515

Before proceeding to the package installation phase, **anaconda** did not check if the *core* package group was available in selected repositories. If this group was not present, the installation terminated unexpectedly. With this update, **anaconda** has been modified to check for the presence of *core*. As a result, a warning message is displayed when *core* is not available, and installation no longer crashes.

BZ#903689

Previously, when configuring a VLAN network device, such as `eth0.171`, during the installation, the same configuration was incorrectly applied also for its parent device. Consequently, the VLAN parent device, such as `eth0`, was incorrectly configured during the installation. The bug has been fixed, and the VLAN device configuration is now applied correctly.

BZ#909463

Under certain circumstances, kernel command-line entries created by **anaconda** and passed to **GRUB** did not work correctly. Consequently, in multi-path configuration, the Boot File System (BFS) terminated unexpectedly when the last FCoE interface specified in kernel command was not on-line. With this update, the form of kernel command-line entries has been modified, and BFS no longer fails in the aforementioned scenario.

BZ#919409

Previously, the `/etc/multipath/bindings` file had incorrect SELinux context after installation. This bug has been fixed, and `/etc/multipath/bindings` is now installed with correct SELinux context.

BZ#921609

Prior to this update, the generated `kickstart` file did not contain correct network commands for VLAN interfaces. Consequently, these commands were not reusable during the installation. This bug has been fixed, and the generated `kickstart` now contains reusable network commands.

BZ#928144

By default, the AMD IOMMU driver is disabled in Red Hat Enterprise Linux 6 for stability reasons. However, when IOMMU is expected to be present for trusted boot, this driver is needed. With this update, **anaconda** has been modified to enable AMD IOMMU in the kernel boot parameters when the `tboot` package is installed. MD IOMMU is enabled when trusted boot is in use and AMD IOMMU specifications are present and enabled in the BIOS. To revert these settings, users may remove the "amd_iommu=on" kernel parameter if stability issues are encountered.

BZ#947704

Previously, it was not possible to blacklist the `usb-storage` module during the installation of Red Hat Enterprise Linux 6. This bug has been fixed, and `usb-storage` can now be blacklisted without complications.

BZ#949409

Under certain rare circumstances, the `dasd_eckd_mod` driver was not loaded during `linuxrc.s390` installation and **anaconda** became unresponsive. With this update, a patch has been applied to prevent this problem.

BZ#971961

Previously, bond network devices were activated only in the early stage of installation. Consequently, bond devices configured by network commands in the `stage2` file were not activated. This behavior has been changed and bond devices can now be activated also in later stages of installation.

BZ#994504

Previously, **anaconda** loaded certain required packages multiple times during installation. Consequently, the dependency solving took a long time, growing with number of disks and file systems. With this update, **anaconda** has been modified to use a more efficient way of selecting packages, thus reducing the time spent on dependency solving.

BZ#998486

With this update, **anaconda** no longer requires the `fcoe-utils` package for installation on the IBM System z architectures.

BZ#1003844

Prior to this update, **anaconda** limited swap size to 10 % of disk space even if **--hibernation** option was used in the **kickstart** file. With this update, **anaconda** has been modified to accept the **--hibernation** option, and swap size is no longer limited to 10% of disk space when this option is specified.

BZ#1004752

Due to an incorrect setting in the **/etc/ssh/sshd_config.anaconda** configuration file, the **sshd** daemon did not start during installation on IBM System z architectures in FIPS mode. Consequently, the installation was not successful. This bug has been fixed, and **sshd** now runs as expected during installation in FIPS mode.

BZ#1007641

Prior to this update, multipath devices were not listed during installation in VNC mode. This bug has been fixed, and these devices are now listed properly.

BZ#1007683

Devices directly formatted with a file system without any partitions are not supported in Red Hat Enterprise Linux 6. Previously, **anaconda** did not verify if devices meet this condition. Consequently, when attempting to create a new partition on such unsupported device, **anaconda** terminated unexpectedly. With this update, **anaconda** has been modified to check if the device is unpartitioned and to abort partitioning in such case, thus preventing the crash.

BZ#1007884

Previously, a bug in the zipl boot loader caused a runtime error in **anaconda**. Consequently, the IBM System z architectures with rootfs on iSCSI LUN failed to boot after an **anaconda** upgrade from Red Hat Enterprise Linux 6.4 to 6.5. This bug has been fixed, and the failed booting no longer occurs after system upgrade.

BZ#1008731

Due to an outdated FCoE detection for Broadcom adapters in **anaconda**, the system was unable to boot after OS after FCoE BFS installation on HP systems. With this update, **anaconda** has been modified to correctly detect FCoE on Broadcom adapters, and the boot problems no longer occur in the aforementioned scenario.

BZ#1008941

Under certain circumstances, after an upgrade from Red Hat Enterprise Linux 6.4 to 6.5, the IBM System z system did not boot from the correct storage device. This bug has been fixed, and System z systems now boot from the correct device after upgrade.

BZ#1009691

Certain adapters, such as *10GBaseT Twin Pond* require longer time to link up. This time often exceeded the timeout limit of the **fipvlan** tool used by the installer. Consequently, adding FCoE targets in the GUI failed by timing out. With this update, the timeout limit of **fipvlan** has been raised. As a result, FCoE target is now added successfully regardless of adapter type. Nevertheless, to view the added device in the GUI, user has to go two screens back to the language selection and then forth.

BZ#1013176

Previously, the list of FCoE LUNs disappeared from the SAN Devices tab in **anaconda** after adding a second adapter during installation of the Specialized Storage BFS. This bug has been fixed, and the list is now displayed correctly during the installation.

BZ#1018703

Prior to this update, **anaconda** incorrectly extracted partition names for NVMe devices. Consequently, the boot loader installation failed on NVMe devices. This bug has been fixed, and NVMe devices are now installed successfully.

Enhancements**BZ#890095**

This update adds more flexible support for disk references within the `--driveorder` option in the kickstart boot loader. It is now possible to specify disks that use the `/dev/disk/by-*/` folders as arguments for `--driveorder`.

BZ#905227

This update adds the `--ipv6gateway` option to the kickstart network command, which allows to specify a default IPv6 gateway. Now, both IPv4 and IPv6 default gateways can be specified in network kickstart command using `--gateway` or `--ipv6gateway` respectively.

BZ#915666

With this update, a partition size check has been added to **anaconda** to ensure that the boot partition on x86 architectures is always less than 2TB, which is required by the **GRUB** boot loader.

BZ#917815

With this update, **anaconda** has been modified to allow the DDNS method in the installer. If a **hostname** is specified in the kickstart configuration of a network device that uses the DHCP protocol, this **hostname** is passed to the **dhclient** utility.

Users of anaconda are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.3. ARPTABLES_JF

8.3.1. RHBA-2013:0843 — arptables_jf bug fix update

Updated arptables_jf packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The arptables_jf utility controls the arpfilter packet filtering code in the Linux kernel.

Bug Fixes**BZ#807315**

Prior to this update, both the "mangle-hw-s" and "mangle-hw-d" options required the use of the "--arhln" option. However, even if the "--arhln" option was specified on the command line, the "arptables" command did not recognize it. As a consequence, it was not possible to use those two options successfully. These updated packages fix this bug and the "--arhln" option can now be used together with the mangle hardware options.

BZ#963209

When the "-x" command line option (exact values) was used along with the "-L" (List rules) option, the arptables utility did not list rules but issued an error message saying "-x" option is illegal with "-L". With this update, the arptables utility now uses the "-x" option when listing rules.

Users of arptables_jf are advised to upgrade to these updated packages, which fix these bugs.

8.4. AUGEAS

8.4.1. RHSA-2013:1537 — Low: augeas security, bug fix, and enhancement update

Updated augeas packages that fix two security issues, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, are available for each vulnerability from the CVE links associated with each description below.

Augeas is a utility for editing configuration. Augeas parses configuration files in their native formats and transforms them into a tree. Configuration changes are made by manipulating this tree and saving it back into native configuration files. Augeas also uses "lenses" as basic building blocks for establishing the mapping from files into the Augeas tree and back.

Security Fix

[CVE-2012-0786](#), [CVE-2012-0787](#)

Multiple flaws were found in the way Augeas handled configuration files when updating them. An application using Augeas to update configuration files in a directory that is writable to by a different user (for example, an application running as root that is updating files in a directory owned by a non-root service user) could have been tricked into overwriting arbitrary files or leaking information via a symbolic link or mount point attack.



NOTE

The augeas package has been upgraded to upstream version 1.0.0, which provides a number of bug fixes and enhancements over the previous version. ([BZ#817753](#))

Bug Fixes

[BZ#799885](#)

Previously, when single quotes were used in an XML attribute, Augeas was unable to parse the file with the XML lens. An upstream patch has been provided ensuring that single quotes are handled as valid characters and parsing no longer fails.

[BZ#855022](#)

Prior to this update, Augeas was unable to set up the "require_ssl_reuse" option in the vsftpd.conf file. The updated patch fixes the vsftpd lens to properly recognize this option, thus fixing this bug.

[BZ#799879](#)

Previously, the XML lens did not support non-Unix line endings. Consequently, Augeas was unable to load any files containing such line endings. The XML lens has been fixed to handle files with CRLF line endings, thus fixing this bug.

BZ#826752

Previously, Augeas was unable to parse `modprobe.conf` files with spaces around "=" characters in option directives. The `modprobe` lens has been updated and parsing no longer fails.

All Augeas users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.5. AUTOFS

8.5.1. RHBA-2013:1690 — autofs bug fix and enhancement update

Updated `autofs` packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The `autofs` utility controls the operation of the **automount** daemon. The daemon automatically mounts file systems when in use and unmounts them when they are not busy.

Bug Fixes

BZ#859078

Under certain circumstances, the **autofs** utility did not respect all configured settings and used the UDP protocol to probe availability of network file systems. This can lead to some servers refusing the connection with the message:

```
Client x.x.x.x is violating the NFSv4 specification by sending a UDP/IP datagram to the NFSv4 server.
```

With this update, **autofs** has been modified to respect explicitly defined NFSv4 requests, thus fixing this bug.

BZ#886623

Due to changes made to the **autofs** utility, when probing server availability at mount time, mounts using the RDMA protocol are no longer recognized. With this update, **autofs** has been modified not to probe availability for mounts that use the RDMA protocol.

BZ#903944

Previously, the **autofs** utility ignored the `--random-multimount-selection` option. Consequently, this setting was not used when mounting local file systems even when it was given. This bug has been fixed and `--random-multimount-selection` now works as expected.

BZ#908020

Previously, when two nearly simultaneous mount requests appeared, NFS mounts mounted by the **autofs** utility sometimes terminated. This was caused by using invalid protoent structures to identify the protocol. With this update, **autofs** has been modified to use numeric protocol IDs, instead of protoent structures. As a result, attempts to mount NFS no longer fail in the described scenario.

BZ#971131

Prior to this update, the **autofs** master map parser did not recognize the **SELinux context=** option and returned a syntax error when the option was used. The master map parser has been updated to recognize **SELinux context=** that can now be used without complications.

BZ#974884

Previously, the **autofs** utility did not recognize the allowed limit of maximum opened files after it was increased by the system administrator. Consequently, the default limit was used regardless of the new configuration. With this update, **autofs** has been modified to check for changes of this limit and to apply them correctly.

BZ#996749

Previously, the **libldap** library was not initialized in a thread-safe manner. Consequently, when running **automount**, the **ber_memalloc_x()** function could have terminated unexpectedly with a segmentation fault. With this update, the initialization of **libldap** has been modified to be thread-safe and **ber_memalloc_x()** no longer crashes in the aforementioned scenario. (BZ#996749)

BZ#979929

When the **automount** daemon was checking host availability and one of the network interfaces was marked "DOWN", **automount** terminated with a segmentation fault. With this update, a check for this case has been added and the segmentation fault no longer occurs.

BZ#994296

When the **automount** daemon received a shutdown signal, executing the **autofs reload** command caused **automount** to stop running when multiple maps were being removed from the auto.master map. A patch has been added to fix this bug and **automount** no longer terminates in the described case.

BZ#994297

A change that removed a code for adding the current map entry caused wildcard indirect multi-mount map entries to fail to mount. A patch to fix wildcard multi-map regression has been added and map entries now mount successfully.

BZ#1002896

Due to an execution order race that occurred when creating an expire thread, the **automount** daemon became unresponsive. The code that handled the expire thread creation has been modified to prevent the aforementioned problem.

BZ#996749

Previously, no locking was performed around LDAP initialization calls. However, these functions are not thread-safe and race conditions could have occurred. With this update, the locking has been added and the risk of race condition is now reduced.

Enhancements**BZ#982103**

The description of the **TIMEOUT** configuration option has been enhanced in the **autofs** man page. The description now explains the internal default configuration more clearly.

BZ#852327

The **autofs** utility has been updated to provide the ability to dump its mount maps in a simple **<key, value>** format in addition to the existing informational format.

Users of autofs are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.6. BATIK

8.6.1. RHBA-2013:1530 — batik bug fix update

Updated batik packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The batik packages provide the Batik toolkit based on a Java technology. This toolkit is used by applications that require to use images in the Scalable Vector Graphics (SVG) format for various purposes, such as viewing, generation, or manipulation.

Bug Fixes

BZ#631677

This update removes the empty batik-debuginfo package.

BZ#867701, BZ#995471

Previously, an attempt to use the rasterizer utility to convert an SVG image to the JPEG format caused an error to be returned. This update applies a patch to fix this bug and rasterizer now converts SVG images to the JPEG format correctly.

BZ#883464

Previously, the manifest.mf file included the keyword "version" instead of "bundle-version". Consequently, the Eclipse platform did not work correctly with Batik utilities. This bug has been fixed and Eclipse now works as expected.

BZ#979527, BZ#995471

Due to a bug in the underlying source code, an attempt to use the ttf2svg font converter failed with an exception. This update applies a patch to fix this bug and ttf2svg now works correctly.

BZ#995471

Previously, the batik packages contained many bugs, among others classpath errors and errors connected with a missing module for handling the JPEG format. Consequently, Batik utilities, such as rasterizer, svgpp, and ttf2svg, failed with exceptions. With this update, the underlying source code has been modified to fix these bugs and the aforementioned utilities now work as expected.

Users of batik are advised to upgrade to these updated packages, which fix these bugs.

8.7. BFA-FIRMWARE

8.7.1. RHBA-2013:1549 — bfa-firmware bug fix and enhancement update

Updated bfa-firmware packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The bfa-firmware package contains the Brocade Fibre Channel Host Bus Adapter (HBA) Firmware to run Brocade Fibre Channel and CNA adapters. This package also supports the Brocade BNA network adapter.



NOTE

The bfa-firmware packages have been upgraded to upstream version 3.2.21-1, which provides a number of bug fixes and enhancements over the previous version. (BZ#[928990](#), BZ#[1007100](#))

All users of bfa-firmware are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.8. BIND-DYNDB-LDAP

8.8.1. RHBA-2013:1636 — bind-dyndb-ldap bug fix update

Updated bind-dyndb-ldap packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The dynamic LDAP back-end is a plug-in for BIND that provides back-end capabilities to LDAP databases. It features support for dynamic updates and internal caching that helps to reduce the load on LDAP servers.

Bug Fixes

BZ#[908780](#)

Previously, the bind-dyndb-ldap plug-in did not handle DNS zones without the "idnsUpdatePolicy" attribute properly, which led to a harmless, but misleading error message:

```
zone serial ([zone serial]) unchanged. zone may fail to transfer to slaves.
```

This message was logged after each zone reload or potentially after each change in the affected DNS zone. The bind-dyndb-ldap plug-in has been fixed, so that it no longer prints any error message if the "idnsUpdatePolicy" attribute is not defined in the DNS zone.

BZ#[921167](#)

Previously, the bind-dyndb-ldap plug-in processed update policies with the "zonesub" match-type incorrectly, which led to the BIND daemon terminating unexpectedly during the processing of the update-policy parameter. The bind-dyndb-ldap plug-in has been fixed to process update-policy with the "zonesub" match-type correctly, and so it no longer crashes in this scenario.

BZ#[923113](#)

The bind-dyndb-ldap plug-in processed settings too early, which led to the BIND daemon terminating unexpectedly with an assertion failure during startup or reload. The bind-dyndb-ldap plug-in has been fixed to process its options later, and so no longer crashes during startup or reload.

BZ#[1010396](#)

Prior to this update, the bind-dyndb-ldap plug-in with the default configuration did not establish

enough connections to LDAP server for the pointer record (PTR) synchronization feature and, consequently, the PTR record synchronization failed. With this update, the default number of connections has been raised to four, and the PTR record synchronization now works as expected.

Users of `bind-dyndb-ldap` are advised to upgrade to these updated packages, which fix these bugs.

8.9. BIOSDEVNAME

8.9.1. RHBA-2013:1638 — biosdevname bug fix and enhancement update

Updated `biosdevname` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `biosdevname` packages contain a `udev` helper utility which provides an optional convention for naming network interfaces; it assigns names to network interfaces based on their physical location. The utility is disabled by default, except for on a limited set of Dell PowerEdge, C Series and Precision Workstation systems.



NOTE

The `biosdevname` packages have been upgraded to upstream version 0.5.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#[947841](#))

Bug Fix

BZ#[1000386](#)

Previously, the `addslot()` function returned the same `"dev->index_in_slot"` value for two or more interfaces. As a consequence, more than one network interfaces could be named `"renameN"`. This update restores the logic used to obtain a port number that existed in `biosdevname` version 0.3.11 and, as a result, all interfaces are named as expected.

Users of `biosdevname` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.10. BOOST

8.10.1. RHBA-2013:1187 — boost bug fix update

Updated `boost` packages that fix one bug are now available.

The `boost` packages provide free peer-reviewed portable C++ source libraries with emphasis on libraries which work well with the C++ Standard Library.

Bug Fix

BZ#[820670](#)

The `Boost` package did not contain the `Boost.Math` shared libraries, which include an inverse of trigonometric functions over complex numbers and `gamma`, `beta` and `erf` special functions, as specified in the Technical Report on C++ Library Extensions. This update adds the `boost-math` sub-package, which includes the symbols corresponding to the mentioned functions.

Users of boost are advised to upgrade to these updated packages, which fix this bug.

8.11. BUSYBOX

8.11.1. RHSA-2013:1732 — Low: busybox security and bug fix update

Updated busybox packages that fix one security issue and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available from the CVE link associated with the description below.

BusyBox provides a single binary that includes versions of a large number of system commands, including a shell. This can be very useful for recovering from certain types of system failures, particularly those involving broken shared libraries.

Security Fix

CVE-2013-1813

It was found that the mdev BusyBox utility could create certain directories within /dev with world-writable permissions. A local unprivileged user could use this flaw to manipulate portions of the /dev directory tree.

Bug Fixes

BZ#820097

Previously, due to a too eager string size optimization on the IBM System z architecture, the "wc" BusyBox command failed after processing standard input with the following error:

```
wc: : No such file or directory
```

This bug was fixed by disabling the string size optimization and the "wc" command works properly on IBM System z architectures.

BZ#859817

Prior to this update, the "mknod" command was unable to create device nodes with a major or minor number larger than 255. Consequently, the kdump utility failed to handle such a device. The underlying source code has been modified, and it is now possible to use the "mknod" command to create device nodes with a major or minor number larger than 255.

BZ#855832

If a network installation from an NFS server was selected, the "mount" command used the UDP protocol by default. If only TCP mounts were supported by the server, this led to a failure of the mount command. As a result, Anaconda could not continue with the installation. This bug is now fixed and NFS mount operations default to the TCP protocol.


All busybox users are advised to upgrade to these updated packages, which contain backported patches to correct these issues.

8.12. CA-CERTIFICATES

8.12.1. RHEA-2013:1596 — ca-certificates enhancement update

Updated ca-certificates packages that add various enhancements are now available for Red Hat Enterprise Linux 6.

The ca-certificates package contains a set of CA certificates chosen by the Mozilla Foundation for use with the Internet Public Key Infrastructure (PKI).



NOTE

The ca-certificates package has been upgraded to upstream version 1.94 as released with NSS version 3.15, which provides an updated set of recent Certificate Authorities according to the Mozilla CA Certificate Policy. Also, the update-ca-trust configuration management tool has been added. (BZ#[973727](#), BZ#[1002646](#))

Enhancement

BZ#[544376](#)

This update provides Shared System Certificate Authority storage, a system-wide trust storage for configuration data, required as an input for certificate trust decisions. This is a functionally compatible replacement for classic Certificate Authority configuration files and for the libnssckbi NSS trust module. This feature must be explicitly enabled by an administrator. Refer to the update-ca-trust man page in the ca-certificates package for a more detailed description of the feature.

Users of ca-certificates are advised to upgrade to these updated packages, which add these enhancements.

8.13. CIFS-UTILS

8.13.1. RHBA-2013:1654 — cifs-utils bug fix update

Updated cifs-utils packages that fix one bug are available for Red Hat Enterprise Linux 6.

The SMB/CIFS protocol is a standard file sharing protocol widely deployed on Microsoft Windows machines. This package contains tools for mounting shares on Linux using the SMB/CIFS protocol. The tools in this package work in conjunction with support in the kernel to allow one to mount a SMB/CIFS share onto a client and use it as if it were a standard Linux file system.

Users of cifs-utils are advised to upgrade to these updated packages, which fix this bug.

8.14. CJKUNI-FONTS

8.14.1. RHBA-2013:0962 — cjkuni-fonts bug fix update

Updated cjkuni-fonts packages that fix one bug are now available.

CJK Unifonts are Unicode TrueType fonts derived from original fonts made available by Arphic Technology under the Arphic Public License and extended by the CJK Unifonts project.

Bug Fix

BZ#651651

Previously, under some configurations, the KDE startup menu did not show any Chinese characters in Chinese locales (both zh-CN and zh-TW), while Japanese and Korean did not have this problem. With this update, the KDE startup menu now displays Chinese characters in Chinese locales.

Users of `ckjuni-fonts` are advised to upgrade to these updated packages, which fix this bug.

8.15. CLUSTER AND GFS2-UTILS

8.15.1. RHBA-2013:1617 — cluster and gfs2-utils bug fix update

Updated cluster and gfs2-utils packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Cluster Manager is a collection of technologies working together to provide data integrity and the ability to maintain application availability in the event of a failure. Using redundant hardware, shared disk storage, power management, and robust cluster communication and application failover mechanisms, a cluster can meet the needs of the enterprise market.

Bug Fixes

BZ#996233

Prior to this update, if one of the `gfs2_tool`, `gfs2_quota`, `gfs2_grow`, or `gfs2_jadd` commands was killed unexpectedly, a temporary GFS2 metadata mount point used by those tools could be left mounted. The mount point was also not registered in the `/etc/mtab` file, and so the `umount -a -t gfs2` command did not unmount it. This mount point could prevent systems from rebooting properly, and cause the kernel to panic in cases where it was manually unmounted after the normal GFS2 mount point. This update corrects the problem by creating an mtab entry for the temporary mount point, which unmounts it before exiting when signals are received.

BZ#893925

Previously, the `cman` utility did not work correctly if there was a brief network failure in a cluster running in `two_node` mode with no fence delay. Consequently, the two nodes killed each other when the connection was re-established. This update adds a 5-second delay to the “fenced” daemon for the node with the higher node ID and the described problem no longer occurs. Another option is to add a fence delay into the `cluster.conf` file, as documented in the Red Hat Knowledgebase (see <https://access.redhat.com/site/solutions/54829>).

BZ#982670

Prior to this update, the `cman` init script did not handle its lock file correctly when executing the `restart` command. Consequently, the node could be removed from the cluster by other members during the node reboot. The `cman` init script has been modified to handle the lock file correctly, and no fencing action is now taken by other nodes of the cluster.

BZ#889564

Previously, when the `corosync` utility detected a “process pause”, an old, therefore invalid, control group ID was occasionally sent to the `gfs_controld` daemon. Consequently, `gfs_controld` became unresponsive. This update fixes `gfs_controld` to discard messages with old control group IDs, and `gfs_controld` no longer hangs in this scenario.

BZ#888857

Prior to this update, the "fenced" daemon and other related daemons occasionally closed a file descriptor that was still referenced by the corosync libraries during an attempt to stop the daemons. Consequently, the daemons did not terminate properly and shutting down the cluster utility failed. This bug has been fixed, the file descriptor now stays open and it is marked unused by the daemons, and the daemons terminate properly.

BZ#989647

Previously, the fsck.gfs2 utility did not handle a certain type of file system corruption properly. As a consequence, fsck.gfs2 terminated with an error message and did not repair the corruption. This update extends the abilities of fsck.gfs2 to handle file system corruption and the described problems no longer occur.

BZ#1007970

Previously, the "-K" option was unavailable in the mkfs.gfs2 utility. Consequently, mkfs.gfs2 returned the "invalid option" error message, and it was impossible to use this option to keep and not to discard unused blocks. With this update, mkfs.gfs2 handles the "-K" option properly.

BZ#896191

The cluster.conf(5) manual page contained incorrect information that the default syslog facility was "daemon". This update corrects this statement to "local4".

BZ#902920

Previously, the fsck.gfs2 utility did not correctly recognize cases when information about a directory in the Global File System 2 (GFS2) was misplaced. Also, fsck.gfs2 did not properly check consistency of the GFS2 directory hash table. As a consequence, fsck.gfs2 did not report problems with the file system and the files in the corrupted directories were unusable. With this update, fsck.gfs2 has been modified to do extensive sanity checking and it is now able to identify and fix the described problems among others.

BZ#963657

Prior to this update, nested Global File System 2 (GFS2) mount points were not taken into account when stopping the GFS2 resources. Consequently, the mount points were not being unmounted in the correct order and the gfs2 utility failed to stop. The gfs2 init script has been modified to unmount GFS2 mount points in the correct order and the stopping of gfs2 no longer fails in this scenario.

BZ#920358

Previously, the qdiskd daemon did not correctly handle newly rejoined nodes that had been rebooted uncleanly. Consequently, qdiskd removed such nodes after its initialization. With this update, qdiskd skips counting of the missed updates for nodes in the "S_NONE" state, and it no longer removes nodes in the described scenario.

BZ#888318

Previously, the qdiskd daemon did not issue a specific error message for cases when the token timeout was set incorrectly in the "cluster.conf" file. Consequently, qdiskd terminated with the "qdiskd: configuration failed" error message giving no details. This update adds a specific error message for the described cases.

BZ#886585

Previously, the gfs2_grow utility returned a zero exit status even in cases where no growth was possible, due to how little the device had grown. Consequently, automated scripts, used especially for

testing of `gfs2_grow`, received an incorrect "0" return code. With this update, `gfs2_grow` has been modified to return a non-zero exit status when its operations fail.

BZ#871603

Previously, the help text for the `"ccs_tool create"` command contained incorrect parameters for the `"addfence"` subcommand, namely `"user"` instead of `"login"`. Consequently, users could create an incorrect `"cluster.conf"` file. With this update, the help text has been corrected.

BZ#985796

Previously, when the `fsck.gfs2` utility was repairing the superblock, it looked up the locking configuration fields from the `"cluster.conf"` file. Consequently, the `"lockproto"` and `"locktable"` fields could be set improperly when the superblock was repaired. With this update, the `"lockproto"` and `"locktable"` fields are now set to sensible default values and the user is now instructed to set the fields with the `tunegfs2` utility at the end of the `fsck.gfs2` run.

BZ#984085

Previously, the `fsck.gfs2` utility did not properly handle cases when directory leaf blocks were duplicated. As a consequence, files in the corrupted directories were occasionally not found and `fsck.gfs2` became unresponsive. With this update, `fsck.gfs2` checks for duplicate blocks in all directories, identifies and fixes corruptions, and it no longer hangs in this scenario.

Users of `cluster` and `gfs2-utils` are advised to upgrade to these updated packages, which fix these bugs.

8.16. CLUSTERMON

8.16.1. RHBA-2013:1602 — clustermon bug fix update

Updated `clustermon` packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The `clustermon` packages are used for remote cluster management. The `modclusterd` service provides an abstraction of cluster status used by `conga` and by the Simple Network Management (SNMP) and Common Information Model (CIM) modules of `clustermon`.

Bug Fixes

BZ#951470

Prior to this update, the `modclusterd` service made an improper CMAN API call when attempting to associate the local machine's address with a particular cluster node entry, but with no success. Consequently, `modclusterd` returned log messages every five seconds. In addition, when logging for CMAN was enabled, membership messages included, messages arising from the CMAN API misuse were emitted. Now, the CMAN API call is used properly, which corrects the aforementioned consequences.

BZ#908728

Previously, the `modclusterd` service terminated unexpectedly in IPv4-only environments when stopped due to accessing uninitialized memory only used when IPv6 was available. With this update, `modclusterd` no longer crashes in IPv4-only environments.

BZ#888543

Previously, the SNMP (Simple Network Management Protocol) agent exposing the cluster status and

shipped as cluster-snmp caused the SNMP server (snmpd) to terminate unexpectedly with a segmentation fault when this module was loaded, and the containing server was instructed to reload. This was caused by an improper disposal of the resources facilitated by this server, alarms in particular. Now, the module properly cleans up such resources when being unloaded, preventing the crash on reload.

Users of clustermon are advised to upgrade to these updated packages, which fix these bugs.

8.17. COMPAT-OPENMPI

8.17.1. [RHBA-2013:1711](#) — [compat-openmpi bug fix update](#)

Updated compat-openmpi packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The compat-openmpi packages contain shared libraries from earlier versions of Open Message Passing Interface (Open MPI). The libraries from previous releases have been compiled against the current version of Red Hat Enterprise Linux 6, and the packages enable earlier programs to keep functioning properly.

Bug Fix

BZ#876315

The compat-openmpi packages previously did not ensure compatibility with earlier versions of the Open MPI shared libraries. Consequently, the users failed to run certain applications using Open MPI on Red Hat Enterprise Linux 6.3 and later if those applications were compiled against Open MPI versions used on Red Hat Enterprise Linux 6.2 and earlier. After this update, the compat-openmpi packages now maintain compatibility with earlier versions of Open MPI on Red Hat Enterprise Linux 6.

Users of compat-openmpi are advised to upgrade to these updated packages, which fix this bug.

8.18. CONMAN

8.18.1. [RHBA-2013:1677](#) — [conman bug fix and enhancement update](#)

Updated conman packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

ConMan is a serial console management program designed to support a large number of console devices and simultaneous users. ConMan currently supports local serial devices and remote terminal servers.



NOTE

The conman packages have been upgraded to upstream version 0.2.7, which provides a number of bug fixes and enhancements over the previous version. With this update, support for the ipmiopts directive in the conman.conf configuration file has been included. (BZ#[951698](#))

Bug Fix

BZ#891938

Previously, the length range of timezone strings was not sufficient to process all known timezone codes. As a consequence, the conmand daemon failed to start if the timezone name consisted of five or more characters. The maximum string length has been set to 32, and conmand now always starts as expected.

Users of conman are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.19. COOLKEY

8.19.1. RHBA-2013:1699 — coolkey bug fix and enhancement update

Updated coolkey packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

Coolkey is a smart card support library for the CoolKey, Common Access Card (CAC), and Personal Identity Verification (PIV) smart cards.

Bug Fixes

BZ#806038

In previous versions, coolkey always created a bogus e-gate smart card reader to avoid problems with Network Security Services (NSS) and the PC/SC Lite framework when no smart card reader was available. However, e-gate smart cards are no longer available for smart card authentication, and the NSS and pcsc-lite packages have been updated to handle a situation with no e-gate reader attached. Therefore, this bogus reader in coolkey became unnecessary and could cause problems to some applications under certain circumstances. This update modifies the respective code so that coolkey no longer creates a bogus e-gate smart card.

BZ#906537

With a previous version of coolkey, some signature operations, such as PKINIT, could fail on PIV endpoint cards that support both CAC and PIV interfaces. The underlying coolkey code has been modified so these PIV endpoint cards now works with coolkey as expected.

BZ#991515

The coolkey library registered only with the NSS DBM database, however, NSS now uses also the SQLite database format, which is preferred. This update modifies coolkey to register properly with both NSS databases.

Enhancement

BZ#951272

Support for tokens containing Elliptic Curve Cryptography (ECC) certificates has been added to the coolkey packages so the coolkey library now works with ECC provisioned cards.

Users of coolkey are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.20. COREUTILS

8.20.1. [RHSA-2013:1652](#) — Low: coreutils security, bug fix, and enhancement update

Updated coreutils packages that fix three security issues, several bugs, and add two enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE link(s) associated with each description below.

The coreutils package contains the core GNU utilities. It is a combination of the old GNU fileutils, sh-utils, and textutils packages.

Security Fixes

[CVE-2013-0221](#), [CVE-2013-0222](#), [CVE-2013-0223](#)

It was discovered that the sort, uniq, and join utilities did not properly restrict the use of the `alloca()` function. An attacker could use this flaw to crash those utilities by providing long input strings.

Bug Fixes

[BZ#747592](#)

Previously, due to incorrect propagation of signals from child processes, the return values of the "su" command were incorrect and core dump information was not shown in the parent process. With this update, signal propagation from child processes has been fixed and the return values of the "su" command corrected. As a result, core dump messages from child processes are no longer ignored and the "su" command returns correct exit values.

[BZ#749679](#)

Previously, the su command did not wait for the end of its child processes. As a consequence, the su utility might exit before the child process has finished. This bug has been fixed and now "su" waits for the child process to exit.

[BZ#816708](#)

Previously, when invoked with no user name argument, the "id -G" and "id --groups" commands printed the default group ID listed in the password database. Occasionally, this ID was incorrect or not effective, especially when it has been changed. After this update, the aforementioned commands print only effective and real IDs when no user is specified.

[BZ#827199](#)

The "tail -f" command uses inotify for tracking changes in files. For remote file systems [-,.] inotify is not available. In the case of unknown file systems, for example panasas, "tail -f" failed instead of falling back to polling. Now, the list of known file systems is updated and "tail -f" is modified to fall back into polling for unknown file systems. As result, "tail -f" now works correctly, even on unknown file systems, with only a warning about the unknown file system and a fall back to polling.

[BZ#842040](#)

Previously, the "df" command interpreted control characters in the output mount name. As a consequence, it could be inconvenient to read and problematic for scripts when there are control

characters such as "\n" in the output. Problematic characters have been replaced by a question mark sign ("?"), and such output is no longer hard to read.

BZ#867984

Previously, a Red Hat specific patch for multibytes locales support in the core utilities was missing the handling of the "--output-delimiter" option of the "cut" command. As a consequence, the option was ignored if specified. Support for the "--output-delimiter" option has been implemented in coreutils and users can now use this option with multibyte locales.

BZ#889531

Previously, when an "su" session was terminated by a signal, it returned an incorrect exit status. This caused various issues, such as a ksh lockup, to occur. This update fixes the exit status handling and the aforementioned situation no longer occurs.

BZ#911206

Previously, the stat utility used the setpwent() and setgrent() functions. This caused NIS database download problems when the time stat utility was called, thus causing performance issues. After this update, the aforementioned system calls are no longer present in the stat utility source code. As a result, NIS database downloads are not necessary with every stat utility run.

BZ#956143

When parsing a file's content, in which the end of a field was specified using the obsolete key formats (+POS -POS), the sort utility determined the end of the field incorrectly, and therefore produced incorrect output. This update fixes the parsing logic to match the usage of the "-k" option when using these obsolete key formats. The sort utility now returns expected results in this situation.

BZ#960160

Previously, in some cases, the date utility could parse invalid input. This was due to a sign-extending of "other" bytes in the parsing mechanism. This caused unexpected results of some invalid input. The parsing mechanism has been fixed, and, the date utility now correctly recognizes invalid input where appropriate.

BZ#965654

Previously, the "dd" utility produced the transfer statistics output even if the "status=noxfer" was specified. To fix this bug, a new option, "status=none", has been implemented to suppress all informational output. As a result, unnecessary information produced by dd is no longer displayed with this option.

BZ#967623

The "su" utility has a "-p" option, which preserves some of the environmental variables. However, the su(1) manual page incorrectly stated that the whole environment was preserved. After this update, the manual page has been adjusted to list all the preserved environmental variables.

BZ#980061

When moving directories between two file systems, the "mv" utility failed to overwrite an empty directory, which was a violation of the POSIX standard. After this update, mv no longer fails to overwrite an empty destination directory and the POSIX standard rules are obeyed.

BZ#997537

Previously, the "pr" utility used a suboptimal code routine when the "-n" option was specified, and

inconsistent padding with either zeros or spaces. As a consequence, `pr` terminated unexpectedly when the `-n` option was used with a value of 32 or higher. Moreover, the inconsistent padding was hard to parse by scripts. After this update, line numbers are consistently padded by spaces and the program has been improved to handle high values of the `-n` option correctly. As a result, the `pr` utility no longer terminates unexpectedly.

BZ#1006221

Previously, the `tail -f` command did not monitor dead symbolic links properly. As a result, `tail -f` ignored updates to the referent of a symbolic link after the symbolic link was killed. This bug has now been fixed and `tail -f` now notices when the dead symbolic link is revived and resumes tailing the contents of the referent.

Enhancements**BZ#836557**

Before this update, a directory cycle induced by a bind mount was treated as a fatal error, for example a probable disk corruption. However, such cycles are relatively common and can be detected efficiently. The `du` command has been modified to display a descriptive warning and also to return the appropriate non-zero exit value. This allows bind mounts of various services to be handles correctly.

BZ#908980

In Red Hat Enterprise Linux 6, the `dd` command has a `conv` option, which supports various conversion types. This updates adds support for the `sparse` conversion option, used for sparse files. This feature is useful when copying block devices to files to minimize the actual amount of data occupied. In addition, it can be used for managing virtual machine images in different storage types, including iSCSI and NFS.

Users of `coreutils` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.21. COROSYNC

8.21.1. [RHBA-2013:1531 — corosync bug fix and enhancement update](#)

Updated `corosync` packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The `corosync` packages provide the Corosync Cluster Engine and C Application Programming Interfaces (APIs) for Red Hat Enterprise Linux cluster software.

Bug Fixes**BZ#854216**

When running `corosync` on a faulty network with the `failed_to_rcv` configuration option set, `corosync` was very often terminated with a segmentation fault after a cluster node was marked as "failed to receive". This happened because an assert condition was met during a cluster node membership determination. To fix this problem, the underlying code has been modified to ignore the assert if it was triggered by nodes marked as "failed to receive". This is safe because a single node membership is always established in this situation.

BZ#877349

The corosync-notifyd service was not started right after installation because the default configuration of the corosync notifier did not exist. This fix adds the default configuration for this service in the `/etc/sysconfig/corosync-notifyd` file so that corosync-notifyd can now be started right after installation without any additional configuration.

BZ#880598

Due to a bug in the underlying code, the corosync API could read uninitialized memory, and thus return incorrect values when incrementing or decrementing value of certain objects in the configuration and statistics database. This update modifies the respective code to only read 16 bits of memory instead of 32 bits when returning the `[u]int16` type values. The corosync API no longer read uninitialized memory and return correct values.

BZ#881729

Due to a rare race condition in the corosync logging system, corosync could terminate with a segmentation fault after an attempt to dereference a NULL pointer. A pthread mutex lock has been added to a respective formatting variable so that the race condition between log-formatting and log-printing functions is now avoided.

BZ#906432

Previously, corosync did not support IPv6 double colon notation and did not handle correctly closing braces when parsing the `corosync.conf` file. As a consequence, the totem service failed to start when using IPv6. If the configuration file contained additional closing braces, no error was displayed to inform users why was the configuration file not parsed successfully. This update fixes these parsing bugs so the totem service can now be successfully started, and an error message is displayed if the `corosync.conf` file contains additional closing braces.

BZ#907894

Due to multiple bugs in the corosync code, either duplicate or no messages were delivered to applications if the corosync service was terminated on multiple cluster nodes. This update applies a series of patches correcting these bugs so that corosync no longer loses or duplicates messages in this scenario.

BZ#915490

The corosync-fplay utility could terminate with a segmentation fault or result in unpredictable behavior if the corosync fdata file became corrupted. With this update, corosync-fplay has been modified to detect loops in code and properly validate fdata files. To avoid another cause of fdata corruption, corosync now also prohibits its child processes from logging. As a result of these changes, corosync no longer crashes or becomes unresponsive in this situation.

BZ#915769

If a service section in the `corosync.conf` file did not contain a service name, corosync either terminated with a segmentation fault or refused to start an unknown service. With this update, corosync now properly verifies the name key and if no service name is found, returns an error message and exits gracefully.

BZ#916227

The corosync service did not correctly handle a situation when it received an exit request (the SIGINT signal) before the service initialization was complete. As a consequence, corosync became unresponsive and ignored all signals, except for SIGKILL. This update adds a semaphore to ensure that corosync exits gracefully in this situation.

BZ#922671

When running applications that used the Corosync inter-process communication (IPC) library, some messages in the `dispatch()` function were lost or duplicated. With this update, corosync properly verifies return values of the `dispatch_put()` function, returns the correct remaining bytes in the IPC ring buffer, and ensures that the IPC client is correctly informed about the real number of messages in the ring buffer. Messages in the `dispatch()` function are no longer lost or duplicated.

BZ#924261

Sometimes, when an attempt to shut down the corosync service using the "`corosync-cfgtool -H`" command failed and returned the `CS_ERR_TRY_AGAIN` error code, subsequent shutdown attempts always failed with the `CS_ERR_EXISTS` error. The `corosync-cfgtool` utility has been modified to automatically retry the shutdown command, and the Corosync's Cfg library now allows processing of multiple subsequent shutdown calls. The "`corosync-cfgtool -H`" command now works as expected even on heavily loaded cluster nodes.

BZ#947936

If the `uidgid` section of the `corosync.conf` file contained a non-existing user or group, corosync did not display any error. The underlying code has been modified so that corosync now properly verifies values returned by the `getpwnam_r` system call, and displays an appropriate error message in this situation.

BZ#959184

If an IPC client exited in a specific time frame of the connection handshake, the corosync main process received the `SIGPIPE` signal and terminated. With this update, the `SIGPIPE` signal is now correctly handled by the `sendto()` function and the corosync main process no longer terminates in this situation.

BZ#959189

The corosync process could become unresponsive upon exit, by sending the `SIGINT` signal or using the `corosync-cfgtool` utility, if it had open a large number of `confdb` IPC connections. This update modifies the corosync code to ensure that all IPC connection to the configuration and statistics database are closed upon corosync exit so that corosync exits as expected.

Enhancements**BZ#949491**

The corosync daemon now detects when the corosync main process was not scheduled for a long time and sends a relevant message to the system log.

BZ#956739

In order to improve process of problem detection, output of the `corosync-blackbox` command now contains time stamps of events. This feature is backward-compatible so that output (`fdata`) from old versions of corosync is processed correctly.

Users of corosync are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.22. CPUPOWERUTILS

8.22.1. RHBA-2013:1533 — cpupowerutils bug fix and enhancement update

Updated cpupowerutils packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The cpupowerutils packages provide a suite of tools to manage power states on appropriately enabled central processing units (CPU).

Bug Fixes

BZ#886225

Previously, some of the commands in the cpupowerutils packages were missing manual pages. Manual pages for the turbostat, x86_energy_perf_policy, cpufreq-bench, and cpufreq-bench_plot.sh commands have been added, thus fixing this bug.

BZ#886226

If a non-root user tried to run the cpufreq-bench utility, it terminated unexpectedly with a segmentation fault, and an ABRT notification appeared on the desktop. With this update, a warning message is displayed to the user instead, informing them that it is necessary to run the utility as root.

BZ#886227

Prior to this update, the x86_energy_perf_policy utility failed when it tried to open the /dev/cpu/*/msr/ directory. Consequently, a "permission denied" error message was returned. With this update, a new error message explains that the command needs root privileges and x86_energy_perf_policy cleanly exits.

BZ#886228

Previously, the interactive help for the x86_energy_perf_policy utility was short and confusing. The help text has been expanded to clarify the meaning of the command-line options.

BZ#914623

Due to the missing implementation for the "cpupower set -m" command, the error message is returned upon launching the command. Previously, this message wrongly implied that the sched-mc utility is not supported on the system. This update clarifies the message to clearly state that sched-mc is not yet implemented.

BZ#914787

Previously, running the "cpupower -v" or "cpupower --version" commands returned incorrect version information. This bug has been fixed and a selected component of cpupower now reports the correct version-release number.

Enhancement

BZ#852831

Intel turbostat v3.0 utility has been included in Red Hat Enterprise Linux. The utility is used to read current CPU core frequency and active C-states.

Users of cpupowerutils are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.23. CRASH

8.23.1. RHEA-2013:1565 — crash enhancement update

Updated crash packages that add various enhancements are now available for Red Hat Enterprise Linux 6.

The crash packages provide a self-contained tool that can be used to investigate live systems and kernel core dumps created from the netdump, diskdump, kdump, and Xen/KVM "virsh dump" facilities from Red Hat Enterprise Linux.

Enhancements

BZ#902141

Currently, dump files created by the makedumpfile utility using the snappy compression format are now readable by the crash utility. The snappy format is suitable for the crash dump mechanism that requires stable performance in any situation with enterprise application use.

BZ#902144

With this update, dump files created by the makedumpfile utility using the LZO compression format are now readable by the crash utility. The LZO compression format is fast and stable for randomized data.

BZ#1006622

This update adds support for compressed dump files created by the makedumpfile utility that were generated on systems with physical memory requiring more than 44 bits.

BZ#1017930

This update fixes faulty panic-task backtraces generated by the bt command in KVM guest dump files. The bt command now shows a trace when the guest operating system is panicking.

BZ#1019483

This update fixes the CPU number display on systems with 255 or more CPUs during the initialization, by the set command, the ps command, and by all commands that display the per-task header consisting of the task address, PID, CPU and command name. Without the patch, for CPU 255, the sys command displays "NO_PROC_ID", and the other commands show a "-" for the CPU number; for CPU numbers greater than 255, garbage values would be displayed in the CPU number field.

Users of crash are advised to upgrade to these updated packages, which add these enhancements.

8.24. CRASH-GCORE-COMMAND

8.24.1. RHBA-2013:1720 — crash-gcore-command bug fix

Updated crash-gcore-command packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `crash-gcore-command` packages contain an extension module for the `crash` utility that adds a `gcore` command which can create a core dump file of a user-space task that was running in a kernel dumpfile.

Bug Fix

BZ#890232

Due to a backported `madvise/MADV_DONTDUMP` change in the Red Hat Enterprise Linux 6 kernel, VDSO (Virtual Dynamically linked Shared Objects) and `vsyscall` pages were missing in the generated process core dump. With this update, VDSO and `vsyscall` pages are always contained in the generated process core dump.

Users of `crash-gcore-command` are advised to upgrade to these updated packages, which fix this bug.

8.25. CREATEREPO

8.25.1. RHBA-2013:0879 — createrepo bug fix update

An updated `createrepo` package that fixes two bugs is now available for Red Hat Enterprise Linux 6.

The `createrepo` package contains a utility that generates a common metadata repository from a directory of RPM packages.

Bug Fixes

BZ#877301

Previously, a time-stamp check did not pass if a file did not exist. As a consequence, an empty repository was incorrectly flagged as being up to date and the `createrepo --checkts` command performed no action on an empty repository. With this update, missing file is now considered as a failure, and not a pass. The `createrepo --checkts` command now properly creates a new repository when called on an empty repository.

BZ#892657

The `--basedir`, `--retain-old-md`, and `--update-md-path` options were reported only in the `createrepo` utility help message but not in the man page. This update amends the man page and the options are now properly documented in both the help message and the man page.

Users of `createrepo` are advised to upgrade to this updated package, which fixes these bugs.

8.26. CRONIE

8.26.1. RHBA-2013:1681 — cronie bug fix and enhancement update

Updated `cronie` packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

`Cronie` contains the standard UNIX daemon `crond` that runs specified programs at scheduled times and related tools. It is a fork of the original `vixie-cron` and has security and configuration enhancements like the ability to use `pam` and `SELinux`.

Bug Fixes

BZ#697485

Previously, the crond daemon did not drop data about user privileges before calling the popen() system function. Consequently, warnings about changing privileges were written to the /var/log/crond file when the function was invoked by the non-root user. With this update, crond has been modified to drop user privileges before calling popen(). As a result, warnings are no longer logged in this scenario.

BZ#706979

With this update, file permissions of cron configuration files have been changed to be readable only by the root user.

BZ#733697

Prior to this update, the definition of restart in the cron init file was incorrect. Consequently, a failure was incorrectly reported when restarting the crond daemon. The init file has been fixed and the redundant failure message is no longer displayed after crond restart.

BZ#738232

Cron jobs of users with home directories mounted on a Lightweight Directory Access Protocol (LDAP) server or Network File System (NFS) were often refused because jobs were marked as orphaned (typically due to a temporary NSS lookup failure, when NIS and LDAP servers were unreachable). With this update, a database of orphans is created, and cron jobs are performed as expected.

BZ#743473

With this update, obsolete comments have been removed from the /etc/cron.hourly/0anacron configuration file.

BZ#821046, BZ#995089

Due to a bug in cron's support for time zones, planned jobs were executed multiple times. Effects of this bug were visible only during the spring change of time. This bug has been fixed and jobs are now executed correctly during the time change.

BZ#887859

With this update, an incorrect example showing the anacron table setup has been fixed in the anacrontab man page.

BZ#919440

Previously, the crond daemon did not check for existing locks for daemon. Consequently, multiple instances of crond could run simultaneously. The locking mechanism has been updated and running multiple instances of cron at once is no longer possible.

BZ#985888

Prior to this update, the \$LANG setting was not read by the crond daemon. Consequently, cron jobs were not run with the system-wide \$LANG setting. This bug has been fixed and \$LANG is now used by cron jobs as expected.

BZ#985893

Previously, the `crond` daemon used the `putenv` system call, which could have caused `crond` to terminate unexpectedly with a segmentation fault. With this update, `putenv()` has been replaced with the `setenv()` system call, thus preventing the segmentation fault.

BZ#990710

Prior to this update, the `PATH` variable could be set by `cron` or in `crontable`, but could not be changed by a PAM setting. With this update, `PATH` can be altered by PAM setting. As a result, `PATH` can now be inherited from the environment if the `"-P"` option is used.

BZ#1006869

Previously, an incorrect error code was returned when non-root user tried to restart the `crond` daemon. With this update, a correct code is returned in the described case.

Enhancements**BZ#829910**

This update adds the `RANDOM_DELAY` variable that allows delaying job startups by random amount of minutes with upper limit specified by the variable. The random scaling factor is determined during the `crond` daemon startup so it remains constant for the whole run time of the daemon.

BZ#922829

With this update, the `CRON_CORRECT_MAIL_HEADER` environment variable in the `/etc/crond/sysconfig` configuration file has been updated. With this variable enabled, `cron` now sends emails with headers in RFC compliant format.

Users of `crone` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.27. CVS

8.27.1. RHBA-2013:1555 — cvs bug fix and enhancement update

Updated `cvs` packages that fix one bug and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Concurrent Versions System (CVS) is a version control system that can record the history of your files. CVS only stores the differences between versions, instead of every version of every file you have ever created. CVS also keeps a log of who, when, and why changes occurred.

Bug Fix**BZ#671460**

When a CVS client tried to establish a GSSAPI-authenticated connection to a DNS load-balanced cluster node, the authentication failed because each node had a unique host name. With this update, the GSSAPI CVS server has been modified to search for any Kerberos key that matches the `"cvs"` service and any host name. As a result, the CVS server can now authenticate clients using GSSAPI even if the server's host name does not match the domain name, and thus Kerberos principal host name part, common for all cluster nodes. CVS server administrators are advised to deploy two Kerberos principals to each node: a principal matching the node's host name and a principal matching the cluster's domain name.

Enhancement

BZ#684789

Previously, the CVS server did not pass the client address to the Pluggable Authentication Modules (PAM) system. As a consequence, it was not possible to distinguish clients by the network address with the PAM system and the system was not able to utilize the client address for authentication or authorization purposes. With this update, the client network address is passed to the PAM subsystem as a remote host item (PAM_RHOST). Also, the terminal item (PAM_TTY) is set to a dummy value "cvs" because some PAM modules cannot work with an unset value.

Users of cvs are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

8.28. DEVICE-MAPPER-MULTIPATH

8.28.1. RHBA-2013:1574 — device-mapper-multipath bug fix and enhancement update

Updated device-mapper-multipath packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The device-mapper-multipath packages provide tools for managing multipath devices using the device-mapper multipath kernel module.

Bug Fixes

BZ#975676

Device Mapper Multipath (DM-Multipath) did not test pointers for NULL values before dereferencing them in the sysfs functions. Consequently, the multipathd daemon could terminate unexpectedly with a segmentation fault if a multipath device was resized while a path from the multipath device was being removed. With this update, DM-Multipath performs NULL pointer checks in sysfs functions and no longer crashes in the described scenario.

BZ#889429

Prior to this update, the multipathd daemon did not start listening to udev events (uevents) until all the multipath paths that were discovered on system startup had been configured. As a consequence, multipathd was unable to handle paths that were discovered in the meantime. This bug has been fixed and multipathd now handles all paths as expected in the described scenario.

BZ#889441

Due to incorrectly ordered udev rules for multipathd, link priority was not set for multipath paths when creating the multipath device using initramfs udev rules. Consequently, the /dev/disk/by-uuid/<uuid> symbolic links pointed to multipath paths instead of the multipath device. This could lead to boot problems under certain circumstances. With this update, the multipathd udev rules have been ordered correctly so that the aforementioned symbolic links point to the multipath device as expected.

BZ#902585, BZ#994277

Previously, DM-Multipath did not allocate enough space for the sysfs "state" attribute. Consequently, when a path was switched to the "transport-offline" state, a buffer overflow was triggered, resulting in an error message being logged into the system log. Also, DM-Multipath did not handle correctly paths

in the "quiesce" state, which resulted in unnecessary failure of these paths. With this update, DM-Multipath allocates enough space to store all valid values of the sysfs "state" attribute. Paths in the "quiesce" state are now moved to the "pending" state, which prevents the paths from failing.

BZ#928831

Previously, DM-Multipath did not verify whether the kernel supported the "retain_attached_hw_handler" mpath target feature before setting it. Consequently, the multipath devices which had set "retain_attached_hw_handler" did not work on machines with an older kernel without this feature support. With this update, DM-Multipath checks that the kernel supports the "retain_attached_hw_handler" feature before setting it. The multipath devices now work as expected on systems with older kernels utilizing newer versions of DM-Multipath.

BZ#995251

In certain setups, the Redundant Disk Array Controller (RDAC) did not mark a path as down if the target controller reported an asymmetric access state of the target port to be "unavailable". As a consequence, the multipathd daemon repeatedly attempted to send I/O to an unusable path. This bug has been fixed, and multipathd no longer sends I/O to unusable paths in this case.

BZ#1011341

Previously, the kpartx utility did not take into account the actual sector size of the device when creating partitions for the MS-DOS partition table, assuming a fixed size of 512 bytes per sector. Therefore, kpartx created partitions that were 1/8 of the proper size if the device with a sector size of 4 KB used the MS-DOS partition table. With this update, kpartx verifies the device's sector size and calculates the proper partition size if the device uses the MS-DOS partition table.

BZ#892292

When displaying multipath topology for the specified multipath device, DM-Multipath unnecessarily obtained WWIDs for all the multipath paths for all the configured multipath devices. Consequently, the "multipath -l" command took an extensively longer time to complete than expected, especially on systems containing a large number of multipath devices. This behavior has been changed and when displaying topology of the specified multipath devices, the multipath command now acquires WWIDs only for paths belonging to these devices.

BZ#974129

DM-Multipath previously set the `fast_io_fail_tmo` configuration option before setting the `dev_loss_tmo` option. However, a new value of `fast_io_fail_tmo` is not allowed to be greater than or equal to the current value of `dev_loss_tmo`. Therefore, when increasing values of both options and sysfs failed to set `fast_io_fail_tmo` due to the aforementioned limitation, even `dev_loss_tmo` could not have been set to a new value. With this update, if a new value of `fast_io_fail_tmo` would be too high, DM-Multipath sets it to the highest valid value, that is, the current value of `dev_loss_tmo` minus one. When setting both, the `fast_io_fail_tmo` and `dev_loss_tmo` options, `dev_loss_tmo` is now increased first.

BZ#889987

When the `detect_prio` option was set, DM-Multipath did not verify whether a storage device supports asymmetric logical unit access (AULA) before setting up the AULA prioritizer on the device. Consequently, if the device did not support AULA, multipathd failed to detect AULA priority of the paths and emitted an error message to the system log. This bug has been fixed so that DM-Multipath now verifies whether a path can be set with AULA priority before setting up the AULA prioritizer on the storage device.

BZ#875199

Due to a NULL pointer dereference bug, multipathd could terminate with a segmentation fault when removing a failed path to a multipath device. This update adds a NULL pointer test to the code, preventing multipathd from a fail in this scenario.

BZ#904836

When creating partitions for the GUID Partition Table (GPT), the kpartx utility did not account for the actual sector size of the devices with the sector size other than 512 bytes. As a result, kpartx created partitions that did not match the actual device partitions. With this update, kpartx correctly calculates a size of the created partitions to matches the actual block size of the storage device.

BZ#918825

The kpartx utility did not properly release file descriptors allocated for loopback devices, causing file descriptor leaks. This update corrects the kpartx code, and kpartx no longer leaves file descriptors open after releasing loopback devices.

BZ#958091

When the multipath command failed to load a multipath device map with read/write permissions, the multipath device could have been incorrectly set with read-only access. This happened because the multipath command always retried reloading the map table with read-only permissions even though the failure was not caused by an EROFS error. With this update, multipath correctly reloads a multipath device with read-only permissions only if the first load attempt has failed with an EROFS error.

BZ#986767

Previously, DM-Multipath did not prevent creating a multipath device to a tapdev device, which cannot be a subject to multipath I/O due to an unexpected path format. Consequently, if a multipath device was created on top of a tapdev device, multipathd terminated with a segmentation fault on the tapdev device's removal from the system. With this update, tapdev devices are blacklisted by default and this problem can no longer occur.

Enhancements**BZ#947798**

This update adds a new default keyword, "reload_readwrite", to the /etc/multipath.conf file. If set to "yes", multipathd listens to path change events, and if the path has read-write access to the target storage, multipathd reloads it. This allows a multipath device to automatically grant read-write permissions, as soon as all its paths have read-write access to the storage, instead of requiring manual intervention.

BZ#916667

The multipathd daemon now includes major and minor numbers of the target SCSI storage device along with the path's name to messages that are logged upon path's addition and removal. This allows for better association of the path with the particular multipath device.

BZ#920448

In order to keep naming consistency of multipath devices, DM-Multipath now sets the smallest available user-friendly name even when the /etc/multipath/bindings file has been edited manually. If the smallest user-friendly name cannot be determined, DM-Multipath retains previous behavior and sets the multipath device symbolic name to the next available largest name

BZ#924924

A new default parameter, "replace_wwid_whitespace", has been added to the /etc/multipath.conf file. If set to "yes", the scsi_id command in the default configuration section returns WWID with white space characters replaced by underscores for all applying SCSI devices.

Users of device-mapper-multipath are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.29. DEVICE-MAPPER-PERSISTENT-DATA

8.29.1. RHEA-2013:1696 — device-mapper-persistent-data enhancement update

Updated device-mapper-persistent-data packages that add various enhancements are now available for Red Hat Enterprise Linux 6.

The device-mapper-persistent-data packages provide device-mapper thin provisioning (thinp) tools.

Bug Fix

BZ#814790, **BZ#960284**, **BZ#1006059**, **BZ#1019217**

This enhancement update adds important thin provisioning tools (repair, rmap, and metadata_size) as well as caching tools (check, dump, restore, and repair) to the device-mapper-persistent-data packages in Red Hat Enterprise Linux 6.

Users of device-mapper-persistent-data are advised to upgrade to these updated packages, which add these enhancements.

8.30. DHCP

8.30.1. RHBA-2013:1572 — dhcp bug fix update

Updated dhcp packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The Dynamic Host Configuration Protocol (DHCP) is a protocol that allows individual devices on an IP network to get their own network configuration information, including an IP address, a subnet mask, and a broadcast address. The dhcp packages provide a relay agent and ISC DHCP service required to enable and administer DHCP on a network.

Bug Fixes

BZ#996518

Previously, the dhcpd daemon or the dhclient utility terminated unexpectedly with a segmentation fault when starting on an InfiniBand network interface card (NIC) with an alias interface and a shared-network defined. Consequently, dhcpd and dhclient could not be used with an alias interface in a different subnet on InfiniBand NICs. A patch has been applied to address this problem, and neither dhcpd nor dhclient now crash in this scenario.

BZ#902966

Prior to this update, if some of the IPv6 addresses were not in the subnet range declared by subnet6 in the range6 statement, the DHCPv6 server incorrectly offered an address which was not from the client's subnet. The range6 statement parsing code has been fixed to check whether its addresses

belong to the subnet, in which the range6 statement was declared. With this update, the DHCPv6 server now fails to start with an error message if the range6 statement is incorrect.

BZ#863936

Previously, the DHCPv4 relay agent (dhcrelay) terminated unexpectedly with a segmentation fault if dhcrelay received a packet over an interface without any IPv4 address assigned. With this update, dhcrelay checks whether the interface has an address assigned prior to further processing of the received packet, and the relay agent no longer crashes in this scenario.

BZ#952126

Previously, when a DHCPv6 request from a DHCPv6 client came from a random port number, the DHCPv6 server sent the reply back to the source port of the message instead of sending it to UDP port 546, which is standard for IPv6. Consequently, the client got the reply on the incorrect port. The reply handling in the DHCPv6 server code has been fixed, and the server now sends replies to UDP port 546.

BZ#978420

Previously, the dhcpd daemon managed memory allocations incorrectly when manipulating objects via the Object Management API (OMAPI). As a consequence, several memory leaks were identified in dhcpd. With this update, memory allocation management has been fixed, and dhcpd no longer leaks memory in this scenario.

BZ#658855

Prior to this update, when the dhclient utility obtained a lease containing the "next-server" option, dhclient did not expose the option to the dhclient-script environment. Consequently, NetworkManager was not able to use the "next-server" option from the dhclient's lease. This bug has been fixed, dhclient now correctly exposes the "next-server" option and NetworkManager can use the option from the dhclient's lease.

BZ#919221

Previously, the dhcpd server was not able to properly handle parsing of a zone definition which contained two or more key statements. As a consequence, dhcpd returned a misleading error message about an internal inconsistency. The zone statement parsing code has been fixed; the error message reported by dhcpd is now more precise in this scenario, saying that there is a multiple key definition for the zone.

BZ#1001742

Previously, when the dhclient utility was running under IPv6 using multiple interfaces, only the last started instance was configured, while others lost connection after the lease-time had expired. Consequently, the last started instance of dhclient received all the DHCPv6 packets, while the other instances failed to communicate with the server. With this update, dhclient is now bound to a specified interface, and multiple instances of dhclient communicate correctly.

Users of dhcp are advised to upgrade to these updated packages, which fix these bugs.

8.31. DOVECOT

8.31.1. RHBA-2013:1736 — dovecot bug fix update

Updated dovecot packages that fix one bug are now available for Red Hat Enterprise Linux 6.

Dovecot is an IMAP server for Linux and other UNIX-like systems, primarily written with security in mind. It also contains a small POP3 server. It supports e-mail in either the maildir or mbox format. The SQL drivers and authentication plug-ins are provided as sub-packages.

Bug Fix

BZ#1010279

Because of a bug in dovecot's SSL parameters generator, installation of Red Hat Enterprise Linux 6 with FIPS mode enabled could become unresponsive when installing the dovecot package. This problem has been fixed and the installation now completes successfully in the described scenario.

Users of dovecot are advised to upgrade to these updated packages, which fix this bug.

8.32. DRACUT

8.32.1. RHSA-2013:1674 — Moderate: dracut security, bug fix, and enhancement update

Updated dracut packages that fix one security issue, several bugs, and add two enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE link(s) associated with each description below.

The dracut packages include an event-driven initramfs generator infrastructure based on the udev device manager. The virtual file system, initramfs, is loaded together with the kernel at boot time and initializes the system, so it can read and boot from the root partition.

Security Fix

CVE-2012-4453

It was discovered that dracut created initramfs images as world readable. A local user could possibly use this flaw to obtain sensitive information from these files, such as iSCSI authentication passwords, encrypted root file system crypttab passwords, or other information.

This issue was discovered by Peter Jones of the Red Hat Installer Team.

Bug Fixes

BZ#610462

Previously, the mkinitrd utility had no manual page accessible by users. This update adds the mkinitrd(8) manual page.

BZ#720684

Previously, the dracut utility did not call the "lvchange" command with the "--yes" option. Consequently, specification of the original logical volume name (rd_LVM_LV) was required when booting an LVM snapshot. With this update, dracut calls "lvchange" with the "--yes" option and booting LVM snapshots is now more intuitive.

BZ#857048

Prior to this update, the dracut utility copied symbolic links from the system to initramfs without following every redirection. As a consequence, initramfs could contain stale symbolic links, causing the system to boot incorrectly. This bug has been fixed; dracut now correctly copies symbolic link redirections, initramfs contains the same layout as the real system, and boot problems no longer occur in this scenario.

BZ#886194

The dracut utility did not take into account all parameters of the `/etc/crypttab` file when setting up crypto devices. Consequently, options and file names in `/etc/crypttab` had no effect in initramfs. With this update, dracut passes options and file names to the `cryptsetup` tool when setting up crypto devices, and options and files in `/etc/crypttab` are now applied correctly.

BZ#910605

Previously, the dracut utility needed a network configuration on the kernel command line to boot with Internet Small Computer System Interface (iSCSI). Consequently, in cases where no network configuration was needed, it was not possible to boot with iSCSI. Now, dracut starts the iSCSI service regardless of the network configuration parameters on the kernel command line, and the problem described no longer occurs.

BZ#912299

Previously, the dracut utility used the `grep` tool without unsetting the `"GREP_OPTIONS"` environment variable. As a consequence, `grep` did not work correctly because of arbitrary options if the user had set `GREP_OPTIONS` while calling `yum` or running dracut. With this update, dracut now unsets `GREP_OPTIONS` and user settings of this variable no longer affect the correct operation of dracut.

BZ#916144

Prior to this update, the multipath configuration file was always included in the initramfs, even if the root device was not a multipath device. Consequently, the administrator had to update initramfs before rebooting when changing the multipath configuration. The dracut utility has been fixed to include the multipath configuration only if the root device is a multipath device. Additionally, the administrator can split the configuration for the root device which is used in initramfs. Currently, dracut recognizes:

- `/etc/multipath-root.conf`
- `/etc/multipath-root/*`
- `/etc/xdrdevices-root.conf`

These files will be used in initramfs as follows:

- `/etc/multipath.conf`
- `/etc/multipath/*`
- `/etc/xdrdevices.conf`

The administrator can make sure that only the specific multipath configuration for the root device is included in initramfs if he does not want the whole configuration to be copied.

BZ#947729

Previously, when using the Red Hat Enterprise Virtualization Hypervisor packaging of the kernel on a live image, the path to the kernel which needed to be verified during the initial boot did not work correctly. Consequently, the checksum test of the kernel in Federal Information Processing Standard

(FIPS) mode failed, and the system did not boot. With this update, the dracut-fips module also looks for the kernel image in different paths and checks those paths with the checksum file in initramfs. As a result, booting an installation in FIPS mode now checks the correct kernel image and if the checksum is correct, the system continues to boot in FIPS mode.

BZ#960729

The dracut utility did not include the xhci-hcd kernel module in the initramfs image. Consequently, the kernel did not recognize USB 3.0 devices in an early boot stage and the root files system could not be mounted from a USB 3.0 disk. With this update, dracut now includes the xhci-hcd driver in initramfs, and the system is able to boot from USB 3.0 disks.

BZ#1011508

Previously, if the "biosdevname=1" parameter had not been specified on the kernel command line, the dracut utility disabled biosdevname network interface renaming on all machines. Consequently, on Dell machines, interfaces used in initramfs did not have automatic biosdevname names, even though biosdevname interface renaming was active later in the boot process. With this update, dracut only disables biosdevname if the parameter is set to "0". For non-Dell machines, biosdevname now renames interfaces only if "biosdevname=1" is specified on the kernel command line, and Dell machines have biosdevname named interfaces in initramfs.

BZ#1012316

Previously, the time necessary to activate Fibre Channel over Ethernet (FcoE) on a 10GBaseT Twin Pond adapter was too long. As a consequence, the fipvlan utility called by dracut timed out in the process of waiting for the link to come up, and the boot failed. With this update, fipvlan is called with a parameter to wait 30 seconds for the link to come up, and the problem no longer occurs.

BZ#1018377

Previously, when the dracut utility was running the ldd tool, ldd forwarded its output to the cat utility to use the SELinux permissions of cat to display the output. Consequently, if the ldd forwarded the output to cat, and cat forwarded the output further, and the pipe reader exited early, cat received an "EPIPE" signal and reported it to the standard error output. With this update, dracut redirects standard error of ldd calls to the /dev/null file, and the error message of cat is now hidden in this scenario.

Enhancements

BZ#851666

The dracut utility now supports bonding of network interfaces in initramfs. Bonding parameters can be specified on the kernel command line in the following format:

```
bond=<bondname>[:<bondslaves>:[:<options>]]
```

This sets up the <bondname> bonding device on top of <bondslaves>. For more information, run the "modinfo bonding" command.

BZ#1012626

The National Institute of Standards and Technology (NIST) now requires the FIPS module to be defined as a cryptosystem. Therefore, this update adds the /etc/system-fips file marker when the dracut-fips rpm package is installed. It provides a stable file location for FIPS product determination to be used by libraries and applications.

All dracut users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.32.2. RHBA-2013:1747 — dracut bug fix update

Updated dracut packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The dracut packages include an event-driven initramfs generator infrastructure based on the udev device manager. The virtual file system, initramfs, is loaded together with the kernel at boot time and initializes the system, so it can read and boot from the root partition.

Bug Fixes

BZ#1029844

In FIPS mode, the self checking of binaries is only done if the `/etc/system-fips` file is present. Prior to this update, the dracut utility did not copy the `/etc/system-fips` file and some checksum files in the initial ram file system (initramfs). As a consequence, the self check of the tools needed to decrypt a partition was not done and the tools terminated unexpectedly. This bug has been fixed, dracut now copies all the needed files in the initramfs, and systems with encrypted disks can now boot successfully in FIPS mode.

BZ#1029846

When booting in FIPS mode on live ISO images, dracut searched for the checksum file of the kernel image in the wrong place. Consequently, the booting process failed. With this update, the path to the checksum file has been corrected, and live ISO images can now boot in FIPS mode as expected.

Users of dracut are advised to upgrade to these updated packages, which fix these bugs.

8.33. E2FSPROGS

8.33.1. RHBA-2013:1689 — e2fsprogs bug fix update

Updated e2fsprogs packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The e2fsprogs packages provide a number of utilities for creating, checking, modifying, and correcting any inconsistencies in the ext2 file systems.

Bug Fixes

BZ#922847

Previously, the `e2fsck` utility was unable to detect inconsistencies related to overlapping interior or leaf nodes in the extent tree. As a consequence, some of ext4 extent tree corruptions were not detected or repaired by `e2fsck` but they were detected by the kernel at run time. With this update, `e2fsck` is able to detect and repair the described problems as expected.

BZ#994615

Previously, the `e2fsck` utility incorrectly detected uninitialized extents past end of file (EOF) as invalid. Consequently, `e2fsck` identified pre-allocated blocks past EOF as corrupt. This bug has been fixed and `e2fsck` now identifies uninitialized extents past EOF correctly.

BZ#873201

The `resize2fs` utility did not properly handle resizing of an `ext4` file system to a smaller size. As a consequence, files containing many extents could become corrupted if they were moved during the resize process. With this update, `resize2fs` maintains a consistent extent tree when moving files containing many extents, and such files no longer become corrupted in the described scenario.

BZ#974975

Previously, the `resize2fs` utility did not correctly relocate inode and block bitmaps when resizing an `ext4` file system to a smaller size. Consequently, some file systems became corrupted when the bitmaps were not moved within the new file system size. A patch has been provided to address this bug and `resize2fs` now maintains a consistent file system in the described scenario.

BZ#885083

Previously, the `e2fsck` utility failed to store information about file system errors correctly. Consequently, entries in the journal were sometimes not properly propagated to the file system superblock. This bug has been fixed and `e2fsck` now handles all file system errors as expected.

BZ#895679

Previously, the `e2fsck` utility did not clear the error log when processing an `ext4` file system. Consequently, `e2fsck` stored detailed error information in the `ext4` file system superblock and returned it periodically upon mounting. With this update, the error log is cleared when `e2fsck` completes and the redundant error messages are no longer returned.

BZ#927541

Prior to this update, the `filefrag` utility occasionally reported incorrect extent counts. A patch has been applied to address this problem and extents are now counted correctly.

Users of `e2fsprogs` are advised to upgrade to these updated packages, which fix these bugs.

8.34. EFIBOOTMGR

8.34.1. RHBA-2013:1687 — efibootmgr bug fix update

Updated `efibootmgr` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `efibootmgr` utility is responsible for the boot loader installation on Unified Extensible Firmware Interface (UEFI) systems.

Bug Fix

BZ#924892

Previously, when an invalid value was passed to the `"efibootmgr -o"` command, the command did not recognize the problem and passed the incorrect value to other functions. This could have lead to several complications such as commands becoming unresponsive. With this update, `efibootmgr` has been modified to test for invalid input. As a result, an error message is displayed in the aforementioned scenario.

Users of `efibootmgr` are advised to upgrade to these updated packages, which fix this bug.

8.35. EMACS

8.35.1. RHBA-2013:1088 — emacs bug fix update

Updated emacs packages that fix a bug are now available for Red Hat Enterprise Linux 6.

GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a scripting language (elisp), and the capability to read email and news.

Bug Fix

BZ#678225

The Lucida Typewriter and Lucida Console fonts were not usable with Emacs 23.1 in Red Hat Enterprise Linux 6. Consequently, the following error message was displayed in the Messages buffer: "set-face-attribute: Font not available". With this update, no error message is displayed in this scenario and the selected font can be used to display the buffer contents.

Users of emacs are advised to upgrade to these updated packages, which fix this bug.

8.36. ENVIRONMENT-MODULES

8.36.1. RHBA-2013:0844 — environment-modules bug fix update

Updated environment-modules packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The environment-modules packages provide for the dynamic modification of a user's environment using modulefiles. Each modulefile contains the information needed to configure the shell for an application. Once the package is initialized, the environment can be modified on a per-module basis using the module command which interprets modulefiles.

Bug Fixes

BZ#918540

When updating the environment-modules package, changes to the `/usr/share/Modules/init/.modulespath` config file were being silently replaced by upgrades. The file is now set marked as `%config(noreplace)` in the spec file, thus it is preserved between updates.

BZ#929007

The environment scripts of `csh` and `tcsh` used the "test" command without specifying the `PATH` variable. That could have possibly resulted in an unexpected behavior as a user binary called "test" could have been run instead. With this update, the "test" binary is called by its full path. Misbehavior caused by calling a random test binary is no longer possible.

BZ#953198

When updating the environment-modules package, changes to environment scripts in `/etc/profile.d` were not preserved. With this update, those scripts have been marked as configuration scripts, thus they are preserved between updates.

All users of environment-modules are advised to upgrade to these updated packages, which fix these bugs.

8.37. ESC

8.37.1. RHBA-2013:1698 — esc bug fix update

Updated esc packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The esc packages contain the Smart Card Manager GUI, which allows user to manage security smart cards. The primary function of the tool is to enroll smart cards, so that they can be used for common cryptographic operations, such as secure e-mail and website access.

Bug Fixes

BZ#920826

The ESC utility did not start when the latest 17 series release of the XULRunner runtime environment was installed on the system. This update includes necessary changes to ensure that ESC works as expected with the latest version of XULRunner.

BZ#961582

The ESC utility can be started manually or automatically when a card is inserted. Previously, when ESC started automatically, the `~/redhat/` directory was created and granted with the read, write and execute permissions. However, some files within this directory had the permissions to read and write only. This inconsistency has been fixed with this update and the permissions are now set properly in the described scenario.

BZ#981156

Due to a bug in the `esc.desktop` file, an error message has been logged in the `/var/log/messages/` directory. This update applies a patch to fix this bug and the error message is no longer returned.

Users of esc are advised to upgrade to these updated packages, which fix these bugs.

8.38. EVOLUTION

8.38.1. RHSA-2013:1540 — Low: evolution security, bug fix, and enhancement update

Updated evolution packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, are available for each vulnerability from the CVE links associated with each description below.

Evolution is the integrated collection of email, calendaring, contact management, communications, and personal information management (PIM) tools for the GNOME desktop environment.

Security Fix

CVE-2013-4166

A flaw was found in the way Evolution selected GnuPG public keys when encrypting emails. This could result in emails being encrypted with public keys other than the one belonging to the intended recipient.

**NOTE**

The Evolution packages have been upgraded to upstream version 2.32.3, which provides a number of bug fixes and enhancements over the previous version. These changes include implementation of Gnome XDG Config Folders, and support for Exchange Web Services (EWS) protocol to connect to Microsoft Exchange servers. EWS support has been added as a part of the evolution-exchange packages. (BZ#883010, BZ#883014, BZ#883015, BZ#883017, BZ#524917, BZ#524921, BZ#883044)

The gtkhtml3 packages have been upgraded to upstream version 2.32.2, which provides a number of bug fixes and enhancements over the previous version. (BZ#883019)

The libgdata packages have been upgraded to upstream version 0.6.4, which provides a number of bug fixes and enhancements over the previous version. (BZ#883032)

Bug Fix**BZ#665967**

The Exchange Calendar could not fetch the "Free" and "Busy" information for meeting attendees when using Microsoft Exchange 2010 servers, and this information thus could not be displayed. This happened because Microsoft Exchange 2010 servers use more strict rules for "Free" and "Busy" information fetching. With this update, the respective code in the openchange packages has been modified so the "Free" and "Busy" information fetching now complies with the fetching rules on Microsoft Exchange 2010 servers. The "Free" and "Busy" information can now be displayed as expected in the Exchange Calendar.

All Evolution users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements. All running instances of Evolution must be restarted for this update to take effect.

8.39. FCOE-TARGET-UTILS

8.39.1. RHBA-2013:1683 — fcoe-target-utils bug fix update

Updated fcoe-target-utils packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The fcoe-target-utils packages contain a command-line interface for configuring FCoE LUNs (Fibre Channel over Ethernet Logical Unit Numbers) and backstores.

Bug Fixes**BZ#854708**

Due to an error leaving a device marked as in-use, attempts to map a block backstore that had been previously mapped would fail. With this update, mappings of block backstores are properly released, and remapping a block device now succeeds.

BZ#880542

Prior to this update, the kernel terminated unexpectedly when the fcoe-target daemon stopped. A patch has been provided to fix this bug, and the kernel now no longer crashes.

BZ#882121

Previously, the target reported support for sequence-level error recovery erroneously. Consequently, interrupting the connection between the FCoE target and a bnx2fc initiator could cause the initiator to erroneously perform sequence-level error recovery instead of exchange-level error, leading to a failure of all devices attached to the target. This bug has been fixed, and connections with a bnx2fc initiator may now be interrupted without disrupting other devices.

BZ#912210

Prior to this update, there was an error in the python-rtplib library. Consequently, when creating a pscsi (SCSI pass-through) storage object in the targetcli utility, the python-rtplib returned a traceback. The error in the library has been fixed, and pscsi storage objects are now created without errors.

BZ#999902

Since the fcoe-utils command-line interface is required by the fcoe-target-utils packages and is not supported on the s390x architecture, fcoe-target-utils will not work properly on s390x, and thus has been removed.

Users of fcoe-target-utils are advised to upgrade to these updated packages, which fix these bugs.

8.40. FCOE-UTILS

8.40.1. RHBA-2013:1637 — fcoe-utils bug fix and enhancement update

Updated fcoe-utils, libhbalinux, libhbaapi, and lldpad packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The fcoe-utils packages provide Fibre Channel over Ethernet (FCoE) utilities, such as the fcoadm command-line utility for configuring FCoE interfaces, and the fcoemon service to configure DCB Ethernet QOS filters.



NOTE

The libhbalinux packages contain the Host Bus Adapter API (HBAAPI) vendor library which uses standard kernel interfaces to obtain information about Fiber Channel Host Buses (FC HBA) in the system.

The libhbaapi library is the Host Bus Adapter (HBA) API library for Fibre Channel and Storage Area Network (SAN) resources. It contains a unified API that programmers can use to access, query, observe, and modify SAN and Fibre Channel services.

The lldpad packages provide a user-space daemon and a configuration utility for Intel's Link Layer Discovery Protocol (LLDP) agent with Enhanced Ethernet support.

The fcoe-utils packages have been upgraded to upstream version 1.0.28, which provides a number of bug fixes and enhancements over the previous version, including support for the virtual N_Port to virtual N_Port (VN2VN) protocol. Moreover, the fcoeadm utility now supports listing Fibre Channel Forwarder (FCF) and Link Error Status Block (LESB) statistics, and also support for the fcoe_sysfs kernel interface has been added. Additionally, documentation updates, a new website, mailing lists, and various minor bug fixes are included in this rebase. (BZ#[829793](#), BZ#[829797](#))

The libhbalinux packages have been upgraded to upstream version 1.0.16, which provides a number of bug fixes and enhancements over the previous version. Also, the documentation has been updated and it now directs the user to the new mailing lists. (BZ#[829810](#))

The libhbaapi packages have been upgraded to upstream version 2.2.9, which provides a number of enhancements over the previous version. Also, the documentation has been updated and it now directs the user to the new mailing lists. (BZ#[829815](#))

The lldpad packages have been upgraded to upstream version 0.9.46, which provides a number of bug fixes and enhancements over the previous version, including 802.1Qbg edge virtual bridging (EVB) module support. Also, FCoE initialization protocol (FIP) application type-length-value (TLV) parsing support, help on usage of the out-of-memory killer, manual page and documentation enhancements have been included. (BZ#[829816](#), BZ#[893684](#))

Bug Fix

BZ#[903099](#)

Due to a bug in the kernel, destroying an N_Port ID Virtualization (NPIV) port while using an ixgbe adapter, the fcoe service init script could become unresponsive on shutdown. An init script patch has been applied to destroy the associated virtual ports first, and the fcoe service no longer hangs in the described scenario.

Enhancement

BZ#[981062](#)

The readme file has been updated with a note clarifying that the file system automounting feature is enabled in the default installation of Red Hat Enterprise Linux 6.


Users of fcoe-utils, libhbalinux, libhbaapi, and lldpad are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.41. FEBOOTSTRAP

8.41.1. RHBA-2013:1535 — febootstrap bug fix update

Updated febootstrap packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The febootstrap package is used by libguestfs to build a small appliance.



NOTE

The febootstrap package has been upgraded to upstream version 3.21, which provides one bug fix over the previous version.

Bug Fix

BZ#902478

Previously, when using febootstrap-supermin-helper with the "-g" option, the command did not set the supplemental groups properly. As a consequence, some groups from the user running libguestfs leaked into the appliance build process. After this update, supplemental groups are set correctly.

Users of febootstrap are advised to upgrade to these updated packages, which fix this bug.

8.42. FENCE-AGENTS

8.42.1. RHBA-2013:1562 — fence-agents bug fix and enhancement update

Updated fence-agents packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

Red Hat fence-agents are a collection of scripts for handling remote power management for cluster devices. They allow failed or unreachable nodes to be forcibly restarted and removed from the cluster.

Bug Fixes

BZ#872308

Previously, the fence agents documentation did not mention how to use the fence_ipmilan agent for fence device HP iLO 3. This update adds this information to the fence_ipmilan(8) manual page.

BZ#896603

Previously, the fence agent fence_cisco_ucs did not respect the "delay" attribute. This bug has now been fixed and fence_cisco_ucs waits the appropriate amount of time, as expected.

BZ#978325

Previously, the fence agent fence_cisco_ucs did not use a proper timeout during the login process, which could have an impact on a successful login. With this update, this timeout is set properly and can be customized by users through the standard configuration methods.

BZ#978326

Previously, the fence agent `fence_cisco_ucs` failed with a traceback error when the hostname could not be resolved to an IP address. With this update, `fence_cisco_ucs` exits with an appropriate error message.

BZ#978328

Previously, the fence agent `fence_scsi` did not provide the correct metadata for the pacemaker "unfence" operation. With this update, an "unfence" operation can be run only on local node.

BZ#912773, BZ#994186

Previously, the fence agent `fence_scsi` did not respect the "delay" attribute. This bug has been fixed and `fence_scsi` now waits the appropriate amount of time. As a result, nodes in a 2-node cluster can no longer fence each other.

BZ#959490

Previously, when using the `fence_bladecenter` agent with the "--ssh" option, the fence agent required also the "--password" or "--identity-file" options. However, this behavior was not documented. As a consequence, when using `fence_bladecenter` with the "--ssh" option only, `fence_bladecenter` failed with an error message which was too generic. This bug has been fixed and a more specific error message is now displayed if `fence_bladecenter` fails to connect.

BZ#887349

Previously, the `fence_scsi(8)` manual page did not mention the "unfence" operation which is required for `fence_scsi` to properly function in a cluster environment. With this update, a comment with information about "unfence" in cluster environment has been added to the `fence_scsi(8)` manual page.

BZ#902404

Previously, when fencing a Red Hat Enterprise Linux cluster node with the `fence_soap_vmware` fence agent, the agent terminated unexpectedly with a traceback if it was not possible to resolve a hostname of an IP address. With this update, a proper error message is displayed in the described scenario.

BZ#905478

Due to incorrect detection on newline characters during an SSH connection, the `fence_drac5` agent could terminate the connection with a traceback when fencing a Red Hat Enterprise Linux cluster node. Only the first fencing action completed successfully but the status of the node was not checked correctly. Consequently, the fence agent failed to report successful fencing. When the "reboot" operation was called, the node was only powered off. With this update, the newline characters are correctly detected and the fencing works as expected.

BZ#981086

Previously, the description of the `fence_ipmilan` "lanplus" option in the `fence_ipmilan(8)` manual page was incomplete. This update improves the description of the "lanplus" option and includes information on its impact on security.

BZ#1014000

Previously, an insecure temporary directory was used by the VMware fence agent, which could be used by a local attacker to overwrite an arbitrary local file by the victim running fence agent. This update removes a dependency on the `python-suds` library, which is vulnerable to a symbolic link attack (CVE-2013-2217), and the VMware fence agent now uses `mkdtemp` to create a unique temporary directory.

Enhancements

BZ#870269

Previously, users of the HP Integrated Lights-Out (iLO) 4 fence device had to use the fence_ipmilan fence agent. This update adds support for the iLO fence device to the fence-agents packages.

BZ#886614

This update adds support for the firmware for APC power switches, version 5. This update also adds changes to the fence agent command line interface.

Users of fence-agents are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.43. FENCE-VIRT

8.43.1. RHBA-2013:1601 — fence-virt bug fix update

Updated fence-virt packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The fence-virt packages provide a fencing agent for virtual machines as well as a host agent, which processes fencing requests.

Bug Fixes

BZ#883588

The respective gzip files for the fence_virt(8) and fence_xvm(8) manual pages were previously created with executable permissions for everybody, which is incorrect. This has been fixed and these files are now properly created with 644 permissions.

BZ#903172

A bug in the fence_virt fencing agent could cause the agent to fail listing the virtual machines that could have been fenced by the fence_virt daemon using the serial channel within a virtual interface. This happened when the virtual machine had been started or live-migrated after starting the fence_virt daemon on the cluster node. The bug has been fixed and fence_virt now lists virtual machines as expected in this scenario.

Users of fence-virt are advised to upgrade to these updated packages, which fix these bugs. Before applying this update, make sure all previously released errata relevant to your system have been applied.

8.44. FIRSTBOOT

8.44.1. RHBA-2013:1595 — firstboot bug fix update

Updated firstboot packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The firstboot utility runs after system installation and guides the user through a series of steps that allows for easier configuration of the machine.

Bug Fix

BZ#876018

The code handling the response to a two-button dialog prompted the user to click one of the buttons. After clicking the close button or pressing the Escape key, the response was ignored, and the post-installation process continued even after disagreeing to the end-user license agreement (EULA) in Red Hat Enterprise Linux 6. With this update, the code has been modified to close the dialog and stay on the underlying screen. As a result, clicking the close button or pressing the Escape key works as expected.

Users of firstboot are advised to upgrade to these updated packages, which fix this bug.

8.45. FOOMATIC

8.45.1. RHBA-2013:1084 — foomatic bug fix update

Updated foomatic packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

Foomatic is a comprehensive, spooler-independent database of printers, printer drivers, and driver descriptions. The package also includes spooler-independent command line interfaces to manipulate queues and to print files and manipulate print jobs. foomatic-rip is a print filter written in C.

Bug Fixes

BZ#661770

The foomatic package could not be rebuilt due to the RPM package spec file having incorrect locations for Perl files. The installation locations have been fixed and the package can now be rebuilt.

BZ#726385

Under certain circumstances, the foomatic-rip CUPS filter could fail, causing print jobs to pass raw data to the printer without being correctly filtered. This was caused by a missing parameter to a logging function. This programming error has been corrected and foomatic-rip now behaves correctly in the described scenario.

Users of foomatic are advised to upgrade to these updated packages, which fix these bugs.

8.46. FPRINTD

8.46.1. RHBA-2013:1738 — fprintd bug fix update

Updated fprintd packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The fprintd packages contain a D-Bus service to access fingerprint readers.

Bug Fix

BZ#1003940

When the Pluggable Authentication Module (PAM) configuration includes the pam_fprintd module, PAM uses the glib2 functions where the dlclose() function is executed to unload the glib2 libraries. However, this method is not designed for multi-threaded applications. When a PAM operation was made, Directory Server on Red Hat Enterprise Linux 6 terminated unexpectedly during the shutdown phase because it attempted to unload the glib2 destructor, which had been previously unloaded by

the `fprintd` service. This update applies a patch to fix this bug so that `fprintd` no longer unloads `glib2` when `pam_fprintd` closes. As a result, the `glib2` libraries are unloaded when Directory Server is closed and therefore the server shuts down gracefully.

Users of `fprintd` are advised to upgrade to these updated packages, which fix this bug.

8.47. FREEIPMI

8.47.1. RHBA-2013:1640 — freeipmi bug fix and enhancement update

Updated `freeipmi` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The FreeIPMI project provides "Remote-Console" (out-of-band) and "System Management Software" (in-band) based on the Intelligent Platform Management Interface specification.



NOTE

The `freeipmi` packages have been upgraded to upstream version 1.2.1, which provides a number of bug fixes and enhancements over the previous version. Among others, this rebase adds the `ipmiseld` daemon and subpackage, and the Serial Over Lan (SOL) command processing. This update also provides more secure permissions for configuration files that recognize remote password configuration. (BZ#951700)

Bug Fixes

BZ#616846, BZ#715605

Prior to this update, the `ipmidetectd` daemon did not fully validate input command-line parameters. Consequently, `ipmidetectd` terminated unexpectedly with a segmentation fault when parsing invalid command-line options. With this update, `ipmidetectd` validates command-line input properly, and therefore no longer crashes in this case.

BZ#818168

Previously, the `bmc-watchdog` daemon did not create the PID file and did not write the PID number into the file. As a consequence, tools depending on missing PID values did not work correctly. This bug has been fixed, the PID number is now stored in the created PID file and the described problems no longer occur.

Users of `freeipmi` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.48. FTP

8.48.1. RHBA-2013:0845 — ftp bug fix update

Updated `ftp` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `ftp` packages provide the standard UNIX command-line File Transfer Protocol (FTP) client. FTP is a widely used protocol for transferring files over the Internet, and for archiving files.

Bug Fix

BZ#861113

Prior to this update, when the FTP client was used from a shell with elevated permissions (through the `su` or the `sudo` utility), it incorrectly assumed the UID from the original login, instead of the user initiating the client. Consequently, the local home directory was incorrect. With this update, the underlying code has been modified to correctly get the login credentials using the `getpwuid(3)` utility function call. Now, the local home directory is set according to the user running the client.

All users of `ftp` are advised to upgrade to these updated packages, which fix this bug.

8.49. GCC

8.49.1. RHBA-2013:1609 — gcc bug fix and enhancement update

Updated `gcc` packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The `gcc` packages provide compilers for C, C++, Java, Fortran, Objective C, and Ada 95 GNU, as well as related support libraries.

Bug Fixes

BZ#906234

Due to the small local buffer for read tokens, GCC (GNU Compiler Collection) could trigger stack smashing protector when reading digraphs in a program. The buffer has been enlarged, and thus the digraph tokens can be read without harming the memory.

BZ#921758

Previously, GCC could terminate unexpectedly when compiling C++ code that contained a structure with the `"va_list"` member field. The initialization of such a structure has been fixed, and GCC no longer crashes on such code.

BZ#959564

Prior to this update, the `libgcc` utility could terminate unexpectedly when unwinding the stack for a function annotated with `"__attribute__((ms_abi))"`. This bug has been fixed by ignoring unwind data for unknown column numbers and `libgcc` no longer crashes.

BZ#967003

Previously, GCC could terminate unexpectedly when processing debug statements. This bug has been fixed by removing the value bound to the variable in such debug statements, and GCC no longer crashes in the described scenario.

Enhancement

BZ#908025

GCC now supports strings with curly braces and vertical bar inside inline assembler code. That is, `'{'`, `'}'`, and `'|'` can now be prefixed with the `'%'` sign; in that case they are not handled as dialect delimiters, but are passed directly to the assembler instead.

Users of gcc are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.50. GDM

8.50.1. RHBA-2013:1708 — gdm bug fix update

Updated gdm packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The GNOME Display Manager (GDM) provides the graphical login screen, shown shortly after boot up, log out, and when user-switching.

Bug Fixes

BZ#712959

Logging into the system with GNOME installed while having set KDE Display Manager (KDM) as the default display manager could sometimes cause the user-switch applet to abort. Consequently, switching the user was impossible unless the applet was reloaded. The underlying code has been modified to prevent interference of multiple queued loads of user information so the user-switch applet is now more resilient to crash in this scenario.

BZ#759174

GDM previously did not forward X Display Manager Control Protocol (XDMCP) indirect queries to the correct port of the appropriate machine. Consequently, the GDM host chooser did not work correctly and XDMCP connection could not be established. With this update, GDM now uses the correct port when redirecting XCDMCP queries and XDMCP connections can be established with the chosen remote host as expected.

BZ#785775, BZ#865832

Previously, GDM displayed login messages for an insufficiently short period of time so that some users were not able to read the messages. This update increases the duration of the time period for which a message is displayed at login time to a minimum of 3 seconds.

BZ#795920

GDM previously did not consult content of the "`~/.dmrc`" file before reading the cached copy of the `dmrc` file in the "`/var/cache/gdm/$USERNAME/`" directory. This behavior could lead to incorrect or inconsistent users environment settings, such as the default graphical desktop session or language, in environments using network-mounted home directories. This happened because changes to "`~/.dmrc`" had no effect on machines to which the users logged in and out before modifying the "`~/.dmrc`" file. With this update, GDM reads "`~/.dmrc`" before "`/var/cache/gdm/$USERNAME/dmrc`" so that updates to the user's environment configuration can take effect.

BZ#818074

When the user switched to the already active session, GDM attempted to clean up temporary internal resources twice. This resulted in spurious error messages being logged in the system log. The underlying code has been fixed so that GDM now cleans up those resources correctly.

BZ#844004

When the PreSession shell script fails, the user is expected to be denied login to the system. GDM previously ignored PreSession failures so that the users were able to proceed with an unauthorized login to the system. This update corrects this behavior so that GDM now fails the login process upon

the PreSession script failure.

BZ#861114

GDM adjusted the width of the login window in accordance with the length of the authentication message. If an authentication message was very long, the login window became unreasonably wide, resulting in text being displayed out of the visible screen. With this update, long authentication messages are automatically wrapped so the login window retains the expected size, and the message is displayed properly.

BZ#874202

When the user logged out of the system or switched runlevel, the gdm-smartcard-worker extension was terminated unexpectedly with a segmentation fault. This update modifies GDM to ensure that gdm-smartcard-worker is brought down gracefully.

BZ#874707

The GDM default greeter did not set the LANG environment variable in canonical form. Consequently, in mixed environment deployments, such as networks containing Mac OS X machines, the LANG encoding was not correctly recognized by non-Linux systems. This update ensures that GDM sets environment variables are in canonical form.

BZ#953552

The gdm-smartcard-worker extension terminated unexpectedly with a segmentation fault upon startup if the system was started without smart card support. The respective code in gdm-smartcard-worker has been modified so this GDM extension no longer crashes in this scenario.

BZ#977560

When using the smart card authentication method with the "disable_user_list=True" option set, entering an incorrect PIN disabled all further smart card logins until the user successfully logged in using a different authentication method. This update properly resets the dialog window in this situation and allows users to repeat smart card authentication attempts.

BZ#1006947

When booting to runlevel 5 on IBM S/390 systems, GDM emitted a warning messages about not being able to start the X server, which were harmless but could confuse the user. The underlying GDM code has been modified to no longer attempt to start the X server on IBM S/390 systems, and the messages are no longer logged to the system logs.

Users of gdm are advised to upgrade to these updated packages, which fix these bugs. GDM must be restarted for this update to take effect. Rebooting achieves this, but changing the runlevel from 5 to 3 and back to 5 also restarts GDM.

8.51. GEGL

8.51.1. RHBA-2013:1021 — [gegl bug fix update](#)

Updated gegl packages that fix one bug are now available for Red Hat Enterprise Linux 6.

GEGL (Generic Graphics Library) is a graph-based image processing framework.

Bug Fix

BZ#620378

Documentation files were installed executable. As a consequence, testing tools failed due to that configuration. To fix this bug, executable bits were removed from documentation files and testing tools now work as expected in the described scenario.

Users of `gegl` are advised to upgrade to these updated packages, which fix this bug.

8.52. GHOSTSCRIPT

8.52.1. RHBA-2013:1624 — ghostscript bug fix update

Updated ghostscript packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The Ghostscript suite contains utilities for rendering PostScript and PDF documents. Ghostscript translates PostScript code to common, bitmap formats so that the code can be displayed or printed.

Bug Fixes

BZ#893775

Due to a bug in a function that copies CID-keyed Type 2 fonts, document conversion attempts sometimes caused the `ps2pdf` utility to terminate unexpectedly with a segmentation fault. A patch has been provided to address this bug so that the function now copies fonts properly and `ps2pdf` no longer crashes when converting documents.

BZ#916162

Due to lack of support for the `TPGDON` option for JBIG2 encoded regions, some PDF files were not displayed correctly. A patch has been provided to add this support so that PDF files using the `TPGDON` option are now displayed correctly.

BZ#1006165

Previously, some PDF files with incomplete ASCII base-85 encoded images caused the ghostscript utility to terminate with the following error:

```
/syntaxerror in ID
```

The problem occurred when the image ended with `"~"` (tilde) instead of `"~>"` (tilde, right angle bracket) as defined in the PDF specification. Although this is an improper encoding, an upstream patch has been applied, and ghostscript now handles these PDF files without errors.

Users of ghostscript are advised to upgrade to these updated packages, which fix these bugs.

8.53. GLIB2

8.53.1. RHBA-2013:1545 — glib2 bug fix and enhancement update

Updated glib2 packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

GLib is a low-level core library that forms the basis for projects such as GTK+ and GNOME. It provides data structure handling for C, portability wrappers, and interfaces for such runtime functionality as an event loop, threads, dynamic loading, and an object system.

**NOTE**

The glib2 packages have been upgraded to upstream version 2.26.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#[883021](#))

Users of glib2 are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.54. GLIBC

8.54.1. [RHSA-2013:1605](#) — Moderate: glibc security, bug fix, and enhancement update

Updated glibc packages that fix three security issues, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The glibc packages provide the standard C libraries (**libc**), POSIX thread libraries (libpthread), standard math libraries (**libm**), and the Name Server Caching Daemon (**nscd**) used by multiple programs on the system. Without these libraries, the Linux system cannot function correctly.

Security Fixes

[CVE-2013-4332](#)

Multiple integer overflow flaws, leading to heap-based buffer overflows, were found in glibc's memory allocator functions (pvalloc, valloc, and memalign). If an application used such a function, it could cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

[CVE-2013-0242](#)

A flaw was found in the regular expression matching routines that process multibyte character input. If an application utilized the glibc regular expression matching mechanism, an attacker could provide specially-crafted input that, when processed, would cause the application to crash.

[CVE-2013-1914](#)

It was found that getaddrinfo() did not limit the amount of stack memory used during name resolution. An attacker able to make an application resolve an attacker-controlled hostname or IP address could possibly cause the application to exhaust all stack memory and crash.

Bug Fixes

[BZ#1022022](#)

Due to a defect in the initial release of the getaddrinfo() system call in Red Hat enterprise Linux 6.0, AF_INET and AF_INET6 queries resolved from the /etc/hosts file returned queried names as

canonical names. This incorrect behavior is, however, still considered to be the expected behavior. As a result of a recent change in `getaddrinfo()`, `AF_INET6` queries started resolving the canonical names correctly. However, this behavior was unexpected by applications that relied on queries resolved from the `/etc/hosts` file, and these applications could thus fail to operate properly. This update applies a fix ensuring that `AF_INET6` queries resolved from `/etc/hosts` always return the queried name as canonical. Note that DNS lookups are resolved properly and always return the correct canonical names. A proper fix to `AF_INET6` queries resolution from `/etc/hosts` may be applied in future releases; for now, due to a lack of standard, Red Hat suggests the first entry in the `/etc/hosts` file, that applies for the IP address being resolved, to be considered the canonical entry.

BZ#552960

The `pthread_cond_wait()` and `pthread_cond_timedwait()` functions for AMD64, Intel 64, and Intel P6 architectures contained several synchronizations bugs. Consequently, when a multi-threaded program used a priority-inherited mutex to synchronize access to a condition variable, some threads could enter a deadlock situation when they were woken up by the `pthread_cond_signal()` function or canceled. This update fixes these synchronization bugs and a thread deadlock can no longer occur in the described scenario.

BZ#834386

The C library security framework was unable to handle dynamically loaded character conversion routines when loaded at specific virtual addresses. This resulted in an unexpected termination with a segmentation fault when trying to use the dynamically loaded character conversion routine. This update enhances the C library security framework to handle dynamically loaded character conversion routines at any virtual memory address, and crashes no longer occur in the described scenario.

BZ#848748

Due to a defect in the standard C library, the library could allocate unbounded amounts of memory and eventually terminate unexpectedly when processing a corrupted NIS request. With this update, the standard C library has been fixed to limit the size of NIS records to the maximum of 16 MB, and the library no longer crashes in this situation. However, it is possible that some configurations with very large NIS maps may no longer work if those maps exceed the maximum of 16 MB.

BZ#851470

Previously, the `ttyname()` and `ttyname_r()` library calls returned an error if the `proc (/proc/)` file system was not mounted. As a result, certain applications could not properly run in a chroot environment. With this update, if the `ttyname()` and `ttyname_r()` calls cannot read the `/proc/self/fd/` directory, they attempt to obtain the name of the respective terminal from the devices known to the system (the `/dev` and `/dev/pts` directories) rather than immediately return an error. Applications running in a chroot environment now work as expected.

BZ#862094

A defect in the standard C library resulted in an attempt to free memory that was not allocated with the `malloc()` function. Consequently, the dynamic loader could terminate unexpectedly when loading shared libraries that require the dynamic loader to search non-default directories. The dynamic loader has been modified to avoid calling the `free()` routine for memory that was not allocated using `malloc()` and no longer crashes in this situation.

BZ#863384

Due to a defect in the `getaddrinfo()` resolver system call, `getaddrinfo()` could, under certain conditions, return results that were not Fully Qualified Domain Names (FQDN) when FQDN results were requested. Applications using `getaddrinfo()` that expected FQDN results could fail to operate correctly.

The resolver has been fixed to return FQDN results as expected when requesting an FQDN result and the `AI_CANONNAME` flag is set.

BZ#868808

The `backtrace()` function did not print call frames correctly on the AMD64 and Intel 64 architecture if the call stack contained a recursive function call. This update fixes this behavior so `backtrace()` now prints call frames as expected.

BZ#903754

Debug information previously contained the name "fedora" which could lead to confusion and the respective package could be mistaken for a Fedora-specific package. To avoid this confusion, the package build framework has been changed to ensure that the debug information no longer contains the name "fedora."

BZ#919562

A program that opened and used dynamic libraries which used thread-local storage variables may have terminated unexpectedly with a segmentation fault when it was being audited by a module that also used thread-local storage. This update modifies the dynamic linker to detect such a condition, and crashes no longer occur in the described scenario.

BZ#928318

When the `/etc/resolv.conf` file was missing on the system or did not contain any nameserver entries, `getaddrinfo()` failed instead of sending a DNS query to the local DNS server. This bug has been fixed and `getaddrinfo()` now queries the local DNS server in this situation.

BZ#929388

A previous fix to prevent logic errors in various mathematical functions, including `exp()`, `exp2()`, `expf()`, `exp2f()`, `pow()`, `sin()`, `tan()`, and `rint()`, created CPU performance regressions for certain inputs. The performance regressions have been analyzed and the core routines have been optimized to raise CPU performance to expected levels.

BZ#952422

Previously, multi-threaded applications using the `QReadWriteLocks` locking mechanism could experience performance issues under heavy load. This happened due to the ineffectively designed `sysconf()` function that was repeatedly called from the Qt library. This update improves the `glibc` implementation of `sysconf()` by caching the value of the `_SC_NPROCESSORS_ONLN` variable so the system no longer spends extensive amounts of time by parsing the `/stat/proc` file. Performance of the aforementioned applications, as well as applications repetitively requesting the value of `_SC_NPROCESSORS_ONLN`, should significantly improve.

BZ#966775

Improvements to the accuracy of the floating point functions in the math library, which were introduced by the RHBA-2013:0279 advisory, led to a performance decrease for those functions. With this update, the performance loss regressions have been analyzed and a fix has been applied that retains the current accuracy but reduces the performance penalty to acceptable levels.

BZ#966778

If user groups were maintained on an NIS server and queried over the NIS compat interface, queries for user groups containing a large number of users could return an incomplete list of users. This update fixes multiple bugs in the compat interface so that group queries in the described scenario now return correct results.

BZ#970090

Due to a defect in the name service cache daemon (nscd), cached DNS queries returned, under certain conditions, only IPv4 addresses even though the AF_UNSPEC address family was specified and both IPv4 and IPv6 results existed. The defect has been corrected and nscd now correctly returns both IPv4 and IPv6 results in this situation.

BZ#988931

Due to a defect in the dynamic loader, the loader attempted to write to a read-only page in memory while loading a prelinked dynamic application. This resulted in all prelinked applications being terminated unexpectedly during startup. The defect in the dynamic loader has been corrected and prelinked applications no longer crash in this situation.

Enhancements**BZ#629823**

Previous versions of nscd did not cache netgroup queries. The lack of netgroup caching could result in less than optimal performance for users that relied on heavily on netgroup maps in their system configurations. With this update, support for netgroup query caching has been added to nscd. Systems that rely heavily on netgroup maps and use nscd for caching will now have their netgroup queries cached which should improve performance in most configurations.

BZ#663641

Previously, if users wanted to adjust the size of stacks created for new threads, they had to modify the program code. With this update, glibc adds a new `GLIBC_PTHREAD_STACKSIZE` environment variable allowing users to set the desired default thread stack size in bytes. The variable affects the threads created with the `pthread_create()` function and default attributes. The default thread stack size may be slightly larger than the requested size due to memory alignment and certain other factors.

BZ#886968

The dynamic loader now coordinates with GDB to provide an interface that is used to improve the performance of debugging applications with very large lists of loaded libraries.

BZ#905575

The glibc packages now provide four Static Defined Tracing (SDT) probes in the libm libraries for the `pow()` and `exp()` functions. The SDT probes can be used to detect whether the input to the functions causes the routines to execute the multi-precision slow paths. This information can be used to detect performance problems in applications calling the `pow()` and `exp()` functions.

BZ#916986

Support for the `MAP_HUGETLB` and `MAP_STACK` flags have been added for use with the `mmap()` function. Their support is dependant on kernel support and applications calling `mmap()` should always examine the result of the function to determine the result of the call.

BZ#929302

Performance of the `sched_getcpu()` function has been improved by calling the Virtual Dynamic Shared Object (VDSO) implementation of the `getcpu()` system call on the PowerPC architecture.

BZ#970776

The error string for the ESTALE error code has been updated to print "Stale file handle" instead of "Stale NFS file handle", which should prevent confusion over the meaning of the error. The error string has been translated to all supported languages.

All glibc users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.55. GLUSTERFS

8.55.1. RHBA-2013:1641 — glusterfs bug fix and enhancement update

Updated glusterfs packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

Red Hat Storage is software only, scale-out storage that provides flexible and affordable unstructured data storage for the enterprise. GlusterFS, a key building block of Red Hat Storage, is based on a stackable user-space design and can deliver exceptional performance for diverse workloads. GlusterFS aggregates various storage servers over network interconnects into one large, parallel network file system.

Bug Fixes

BZ#998778

Previously, the "errno" value was not set correctly during an API failure. Consequently, applications using API could behave unpredictably. With this update, the value is set properly during API failures and the applications work as expected.

BZ#998832

Previously, the glusterfs-api library handled all signals that were sent to applications using glusterfs-api. As a consequence, glusterfs-api interpreted incorrectly all the the signals that were not used by this library. With this update, glusterfs-api no longer handles the signals that it does not use so that such signals are now interpreted properly.

BZ#1017014

Previously, the glfs_fini() function did not return NULL, even if the libgfapi library successfully cleaned up all resources. Consequently, an attempt to use the "qemu-img create" command, which used libgfapi, failed. The underlying source code has been modified so that the function returns NULL when the libgfapi cleanup is successful, and the command now works as expected.

Enhancement

BZ#916645

Native Support for GlusterFS in QEMU has been included to glusterfs packages. This support allows native access to GlusterFS volumes using the libgfapi library instead of through a locally mounted FUSE file system. This native approach offers considerable performance improvements.

Users of glusterfs are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.56. GNOME-SCREENSAVER

8.56.1. RHBA-2013:1706 — gnome-screensaver bug fix update

Updated gnome-screensaver packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The gnome-screensaver packages contain the GNOME project's official screen saver program. The screen saver is designed for improved integration with the GNOME desktop, including themeability, language support, and Human Interface Guidelines (HIG) compliance. It also provides screen-locking and fast user-switching from a locked screen.

Bug Fixes

BZ#905935

Previously, when using the virt-manager, virt-viewer, and spice-xpi applications, users were unable to enter the gnome-screensaver password after the screen saver had started. This occurred only when the virtual machine system used the Compiz compositing window manager. After users had released the mouse cursor, then pressed a key to enter the password, the dialog window did not accept any input. This happened due to incorrect assignment of window focus to applications that did not drop their keyboard grab. With this update, window focus is now properly assigned to the correct place, and attempts to enter the gnome-screensaver password no longer fail in the described scenario.

BZ#947671

Prior to this update, the gnome-screensaver utility worked incorrectly when using an X server that does not support the fade-out function. Consequently, gnome-screensaver terminated unexpectedly when trying to fade out the monitor. This bug has been fixed and gnome-screensaver now detects a potential fade-out failure and recovers instead of crashing.

Users of gnome-screensaver are advised to upgrade to these updated packages, which fix these bugs.

8.57. GPXE

8.57.1. RHBA-2013:1628 — gppe bug fix update

Updated gppe packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The gppe packages provide gPXE, an open source Pre-boot Execution Environment (PXE) implementation and bootloader.

Bug Fix

BZ#972671

A DHCP server can be configured to use the Pre-Boot Execution Environment (PXE) to boot virtual machines using the gPXE utility. Previously, PXE boot failed when the next-server details had come from a different DHCP server. This update applies a patch to fix this bug and PXE boot now works as expected in the described scenario.

Users of gppe are advised to upgrade to these updated packages, which fix this bug.

8.58. GREP

8.58.1. RHBA-2013:0977 — grep bug fix update

Updated grep packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The grep utility searches through textual input for lines which contain a match to a specified pattern and then prints the matching lines. GNU grep utilities include grep, egrep and fgrep.

Bug Fixes

BZ#715295

For some regular expressions, the DFA analysis could insert up to double "positions" than there were leaves. Consequently, there were not enough room to insert all the positions and grep could terminate unexpectedly on certain regular expressions. To fix this problem, space allocation has been increased and grep works as expected in the described scenario.

BZ#797934

When a fixed string pattern was empty while the case-insensitive search was active, grep could terminate unexpectedly. With this update, the check for this case has been added to the code and grep works as expected in the described scenario.

BZ#826997

Previously, the code handling case-insensitive searches could alter a string's byte size while converting it to lower case. Consequently, grep could truncate certain output strings. To fix this bug, the grep code has been modified to correctly handle such cases when the byte size gets altered during the conversion to lower case. As a result, case-insensitive searches work correctly and grep no longer truncates its output.

Users of grep are advised to upgrade to these updated packages, which fix these bugs.

8.59. GRUB

8.59.1. RHBA-2013:1649 — grub bug fix and enhancement update

Updated grub packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The grub packages provide GRUB (Grand Unified Boot Loader), a boot loader capable of booting a wide variety of operating systems.

Bug Fixes

BZ#851706

If the title of the GRUB menu entry exceeded the line length of 80 characters, the text showing the remaining time to a boot was inconsistent and thus appeared to be incorrect. The overflowing text was displayed on a new line and the whole text was moved one line down with every passing second. This update splits the text into two lines, and only the second line is rewritten as a boot countdown proceeds so that GRUB behaves correctly for long menu entries.

BZ#854652

When building a new version of grub packages, GRUB did not remove the grub.info file upon the "make clean" command. As a consequence, the grub.info file did not contain the latest changes after applying an update. To fix this problem, the GRUB Makefile has been modified so the grub.info file is now explicitly removed and generated with every package build.

BZ#911715

The GRUB code did not comply with the Unified Extensible Firmware Interface (UEFI) specification and did not disable an EFI platform's watchdog timer as is required by the specification. Consequently, the system was rebooted if the watchdog was not disabled within 5-minutes time frame, which is undesirable behavior. A patch has been applied that disables the EFI watchdog immediately after GRUB is initialized so that EFI systems are no longer restarted unexpectedly.

BZ#916016

When booting a system in QEMU KVM with Open Virtual Machine Firmware (OVMF) BIOS, GRUB was not able to recognize virtio block devices, and the booting process exited to the GRUB shell. This happened because GRUB did not correctly tested paths to EFI devices. The GRUB code now verifies EFI device paths against EFI PCI device paths, and recognizes disk devices as expected in this scenario.

BZ#918824

GRUB did not comply with the UEFI specification when handling the ExitBootServices() EFI function. If ExitBootServices() failed while retrieving a memory map, GRUB exited immediately instead of repeating the attempt. With this update, GRUB retries to obtain a memory map 5 times before exiting, and boot process continues on success.

BZ#922705

When building a 64-bit version of GRUB from a source package, it fails to link executable during the configure phase, unless a 32-bit version of the glibc-static package is installed. No error message was displayed upon GRUB failure in this situation. This has been fixed by setting the grub packages to depend directly on the /usr/lib/libc.a file, which can be provided in different environments. If the file is missing when building the grub packages, an appropriate error message is displayed.

BZ#928938

When installed on a multipath device, GRUB was unreadable and the system was unable to boot. This happened due to a bug in a regular expression used to match devices, and because the grub-install command could not resolve symbolic links to obtain device statistics. This update fixes these problems so that GRUB now boots as expected when installed on a multipath device.

BZ#1008305

When booting in UEFI mode, GRUB previously allocated memory for a pointer to a structure instead allocating memory for the structure. This rendered GRUB to be unable to finish and pass control to the kernel on specific hardware configurations. This update fixes this problem so GRUB now allocates memory for a structure as expected and successfully passes control to the kernel.

BZ#1017296

Previously, GRUB could not be installed on Non-Volatile Memory Express (NVMe) devices because it was unable to parse a device name during the installation process. This update adds a regular expression support for matching NVMe devices, and GRUB can now be successfully installed on these devices.

Enhancements

BZ#848628

GRUB now provides a new menu option "macappend". When "macappend" is used either in the grub.conf file or on the GRUB command line, the "BOOTIF=<MAC_address>" parameter is appended to the kernel command line. This allows specifying a network interface for Anaconda to use during a PXE boot.

Users of grub are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.60. GRUBBY

8.60.1. RHBA-2013:1713 — grubby bug fix update

Updated grubby packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The grubby packages provide grubby, a command-line tool for displaying and editing GRUB (GRand Unified Bootloader) configuration files.

Bug Fixes

BZ#991197

Previously, the grub.conf file was not properly updated after a kernel update with the tboot bootloader. This was due to a bug in the grubby tool which caused it to improperly interpret the grub.conf stanzas that had tboot in them. This update enables grubby to read the HYPERVISOR and HYPERVISOR_ARGS parameters from the /etc/sysconfig/kernel file in order for tboot to perform as intended.

BZ#999908

Prior to this update, yum and anaconda upgrades could have failed with a kernel panic on the AMD64 and Intel 64 architectures due to the RAM disk image not being found. This only happened when tboot was installed, and the kernel "%post" or "%posttrans" scripts were run. This update adds the initramfs disk image to the grub entry, and kernel panic failures no longer occur in the described scenario.

Users of grubby are advised to upgrade to these updated packages, which fix these bugs.

8.61. GTK2

8.61.1. RHBA-2013:1544 — gtk2 and atk bug fix and enhancement update

Updated gtk2 and atk packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The gtk2 packages provide a multi-platform toolkit for creating graphical user interfaces, GIMP Toolkit (GTK+). GTK+ offers a complete set of widgets and is suitable for small projects as well as complete application suites.



NOTE

The ATK library provides a set of interfaces for adding accessibility support to applications and graphical user interface toolkits. By supporting the ATK interfaces, an application or toolkit can be used with tools such as screen readers, magnifiers, and alternative input devices.

The gtk2 packages have been upgraded to upstream version 2.20.1, which provides a number of bug fixes and enhancements over the previous version. (BZ#[883022](#))

The atk packages have been upgraded to upstream version 1.30.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#[883027](#))

Bug Fixes

BZ#[970594](#)

When rendering the text in a combo box, the GTK+ cell renderer always rendered text that was rendered last time as the first item. Consequently, if the previously rendered text did not match any item in the name set, the first item in the "Categories" combo box in the Contacts view could have been rendered as empty, which affected accessibility and automated tests. This update ensures that the cell renderer is now properly updated and renders items for the current combo box call so the aforementioned problem no longer occurs.

BZ#[979049](#)

Due to a bug in the GtkTreeView interface, the expand arrows in a tree view in Evolution stopped functioning after clicking on an icon in the system tray. This update increases robustness of the tree expanding and collapsing code, which fixes this bug.

Users of gtk2 and atk are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.62. HAPROXY

8.62.1. [RHBA-2013:1619](#) — haproxy bug fix and enhancement update

Updated haproxy packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The haproxy packages provide a reliable, high-performance network load balancer for TCP and HTTP-based applications. It is particularly suited for web sites crawling under very high loads while needing persistence or Layer7 processing.



NOTE

The haproxy packages have been upgraded to upstream version 1.4.24, which provides a number of bug fixes and enhancements over the previous version. (BZ#[947987](#))

Bug Fix

BZ#[903303](#)

Previously, the setuid() and setgid() functions did not work properly. As a consequence, the HAProxy load balancer failed to drop supplementary groups correctly after attempting to drop root privileges.

The behavior of the functions has been modified, and HAProxy now drops all supplementary groups as expected.

Enhancement

BZ#921064

With this update, support for TPROXY has been added to the haproxy packages. TPROXY simplifies management tasks of clients behind proxy firewalls. Also, transparent proxying makes the presence of the proxy invisible to the user.


Users of haproxy are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.63. HDPARM

8.63.1. RHBA-2013:1580 — hdparm bug fix and enhancement update

Updated hdparm packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Hdparm is a useful system utility for setting (E)IDE hard drive parameters. For example, hdparm can be used to tweak hard drive performance and to spin down hard drives for power conservation.



NOTE

The hdparm packages have been upgraded to upstream version 9.43, which provides a number of bug fixes and enhancements over the previous version. These enhancements include creating files with the desired size on the ext4 and xfs filesystems, and a possibility to specify the offset for reading operations when measuring timing performance. Other notable enhancements include an ability to obtain and set the "idle3" timeout value of the Western Digital Green (WDG) hard drive, and an ability to obtain and set the Write-Read-Verify feature for hard drives. (BZ#977800)

Bug Fixes

BZ#639623

Previously, the hdparm utility did not assume that some disk information could be unavailable. As a consequence, hdparm could terminate unexpectedly with no useful output. With this update, proper checks for unsuccessful disk queries have been added, and hdparm now terminates with a more detailed error message.

BZ#735887

Prior to this update, the hdparm utility did not assume that some disk information could be unavailable when the user requested information about how much disk space a file occupied. Consequently, hdparm terminated unexpectedly with no useful output in such a scenario. With this update, proper checks for unsuccessful disk queries have been added. As a result, hdparm now terminates with an error message providing detailed information.

BZ#807056

Previously, the hdparm utility retrieved the hard drive identification data in a way that could cause

errors. As a consequence, `hdparm` failed to obtain the data on some occasions and displayed an unhelpful error message. With this update, the respective system call has been replaced with one that is more appropriate and robust. As a result, the hard drive identification data is now successfully obtained and printed in the output.

BZ#862257

When the `hdparm` utility is unable to obtain the necessary geometry information about a hard drive, it attempts to download firmware. Previously, due to incorrect control statements, `hdparm` could terminate unexpectedly with a segmentation fault at such a download attempt. With this update, control statements checking for system call failures have been added. As a result, if `hdparm` cannot operate on a drive, it displays an error message and exits cleanly.

Users of `hdparm` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.64. HSQLDB

8.64.1. RHBA-2013:1614 — hsqldb bug fix update

Updated `hsqldb` packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The `hsqldb` packages provide a relational database management system written in Java. The Hyper Structured Query Language Database (HSQLDB) contains a JDBC driver to support a subset of ANSI-92 SQL.

Bug Fixes

BZ#996152

Previously, the `/etc/sysconfig/hsqldb` file was not marked as "config(noreplace)". Consequently, reinstallation or update of the packages could overwrite changes to the configuration made by the user. With this update, the configuration file has been marked correctly, and modifications to the file are preserved during reinstallation or update.

BZ#962676

Prior to this update, the `hsqldb` database depended on `java` packages of version 1:1.6.0 or later, which are unavailable on some Red Hat Enterprise Linux 6 platforms. As a consequence, installing the `hsqldb` packages failed with an error message. With this update, `java` packages of version 0:1.5.0 or later are required, and the installation of `hsqldb` now proceeds correctly as expected.

Users of `hsqldb` are advised to upgrade to these updated packages, which fix these bugs.

8.65. HWDATA

8.65.1. RHBA-2013:1612 — hwdata bug fix and enhancement update

An updated `hwdata` package that fixes one bug and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The `hwdata` package contains tools for accessing and displaying hardware identification and configuration data.

Bug Fix

BZ#989142

Previously, certain information about the Red Hat Virtio Small Computer System Interface (SCSI) device was missing from the pci.ids database. Consequently, when using the lspci utility, the device name was not shown correctly and the numeric device ID was shown instead. With this update, the pci.ids database has been modified to provide correct information as expected.

Enhancements

BZ#982659

The PCI ID numbers have been updated for the Beta and the Final compose lists.

BZ#739838

With this update, the pci.ids database has been updated with information about AMD FirePro graphic cards.

BZ#948121

With this update, the pci.ids database has been updated with information about the Cisco VIC SR-IOV Virtual Function with the usNIC capability.

All users of hwdata are advised to upgrade to this updated package, which fixes this bug and adds these enhancements.

8.66. HYPERVKVPD

8.66.1. RHBA-2013:1539 — hypervkvpd bug fix update

Updated hypervkvpd packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The hypervkvpd packages contain hypervkvpd, the guest Hyper-V Key-Value Pair (KVP) daemon. Using VMbus, hypervkvpd passes basic information to the host. The information includes guest IP address, fully qualified domain name, operating system name, and operating system release number. An IP injection functionality enables the user to change the IP address of a guest from the host via the hypervkvpd daemon.

Bug Fixes

BZ#920032

Previously, the hypervkvpd service registered to two netlink multicast groups, one of which was used by the cgred service. When hypervkvpd received a netlink message, it was interpreted blindly as its own. As a consequence, hypervkvpd terminated unexpectedly with a segmentation fault. After this update, hypervkvpd now registers only to its own netlink multicast group and verifies the type of the incoming netlink message. Using hypervkvpd when the cgred service is running no longer leads to a segmentation fault.

BZ#962565

Prior to this update, the hypervkvpd init script did not check if Hyper-V driver modules were loaded into the kernel. If hypervkvpd was installed, it started automatically on system boot, even if the system was not running as a guest machine on a Hyper-V hypervisor. Verification has been added to

the `hypervkvpd` init script to determine whether Hyper-V driver modules are loaded into the kernel. As a result, if the modules are not loaded into the kernel, `hypervkvpd` now does not start, but displays a message that proper driver modules are not loaded.

BZ#977861

Previously, `hypervkvpd` was not built with sufficiently secure compiler options, which could, consequently, make the compiled code vulnerable. The `hypervkvpd` daemon has been built with full read-only relocation (RELRO) and position-independent executable (PIE) flags. As a result, the compiled code is more secure and better guarded against possible buffer overflows.

BZ#983851

When using the `Get-VMNetworkAdapter` command to query a virtual machine network adapter, each subnet string has to be separated by a semicolon. Due to a bug in the IPv6 subnet enumeration code, the IPv6 addresses were not listed. A patch has been applied, and the IPv6 subnet enumeration now works as expected.

Users of `hypervkvpd` are advised to upgrade to these updated packages, which fix these bugs. After updating the `hypervkvpd` packages, rebooting all guest machines is recommended, otherwise the Microsoft Windows server with Hyper-V might not be able to get information from these guest machines.

8.67. IBUS-HANGUL

8.67.1. RHBA-2013:1036 — ibus-hangul bug fix update

Updated `ibus-hangul` packages that fix one bug are now available.

The `ibus-hangul` package is a Korean language input engine platform for the IBus input method (IM).

Bug Fix

BZ#965554

Previously, the Hangul engine for IBus did not function properly. If a preedit string was available, and the input focus was moved to another window, then the preedit string was committed. After that, when the input focus was moved back to the window, the X Input Method (XIM) could not handle the first key input. This update resolves this issue with a change in the code, and key press inputs after a focus change are no longer lost in the described scenario.

Users of `ibus-hangul` are advised to upgrade to these updated packages, which fix this bug.

8.68. ICEDTEA-WEB

8.68.1. RHBA-2013:1584 — icedtea-web bug fix and enhancement update

Updated `icedtea-web` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The IcedTea-Web project provides a Java web browser plug-in and an implementation of Java Web Start, which is based on the Netx project. It also contains a configuration tool for managing deployment settings for the plug-in and Web Start implementations.

**NOTE**

The icedtea-web packages have been upgraded to upstream version 1.4.1, which provides a number of bug fixes and enhancements over the previous version including support for updated versions of OpenJDK6 and OpenJDK7. (BZ#[916161](#), BZ#[975098](#))

Users of icedtea-web are advised to upgrade to these updated packages, which fix these bugs and add these enhancements

8.69. INITSCRIPTS

8.69.1. [RHBA-2013:1679](#) — [initscripts bug fix and enhancement update](#)

Updated initscripts packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The initscripts package contains basic system scripts to boot the system, change runlevels, activate and deactivate most network interfaces, and shut the system down cleanly.

Bug Fixes

BZ#[915659](#)

A regular expression, which was used to match the name of the master bond device in the grep utility, was incorrect. Consequently, network scripts did not properly handle lines in interface configuration containing comments and the ifup-eth command failed to activate slave devices. This update provides an updated regular expression for grep and ifup-eth now works as expected in the described scenario.

BZ#[919217](#) BZ#[963944](#)

In Red Hat Enterprise Linux 6.4, a master device was always started after its slaves while using Mode 6 bonding. As a consequence, bonded interfaces were unusable. This update ensures the master device is always set up before its slaves and Mode 6 bonding now works as expected.

BZ#[984003](#)

Previously, mounting of the /proc directory in the initrd script did not take into account options set in the /etc/fstab file. As a consequence, /proc was not mounted with the specified options. With this update, /proc is now re-mounted in the rc.sysinit script, which ensures it is mounted with the specified options.

BZ#[877928](#)

Previously, initscripts called the nmcli utility to stop the interface even if it was not managed by NetworkManager at the time. As a consequence, the interface was stopped, but the output of nmcli stated the action had failed. After this update, nmcli is no longer called when NetworkManager is not handling the interface, for example when it has failed, is disconnected, unmanaged or unavailable. As a result, the output from nmcli now matches the real result.

BZ#[836233](#)

If assigning an IP address through the Dynamic Host Configuration Protocol version 4 (DHCPv4) failed, initscript exited with an error. As a consequence, static IPv4 and IPv6 addresses were not set if DHCPv4 failed. The option IPV4_FAILURE_FATAL has been added to let the user decide whether

the script should continue or exit when DHCPv4 fails. Additionally, if set to "no" and DHCPv6 is enabled in the configuration file, initscript tries to get an IPv6 address even if DHCPv4 fails.

BZ#843402

After sending the TERM signal, the killproc() function always waited \$delay seconds before it checked the process again. This waiting was unnecessary and with this update killproc() checks multiple times during the waiting delay. As a result, killproc can continue almost immediately after a process ends.

BZ#864802

Previously, initscript did not follow the order of mounts specified by the administrator, because some mount types were prioritized. As a consequence, a subdirectory could be mounted before its parent directory. After this update, NFS, the Common Internet File System (CIFS), the Server Message Block (SBM) and other mount types are the last to be mounted. As a result, the mounts in the /etc/fstab file are processed in the right order.

BZ#814427

Previously, the securetty utility always tried to open the /etc/securetty file in read and write mode. As a consequence, on a read-only root filesystem, this led to failure and the file was not modified even if the TTY had already existed. With this update, securetty now checks whether the /etc/securetty file needs to be modified and exits if it does not. As a result, securetty now works correctly on a read-only root filesystem.

BZ#948824

Prior to this update, users were not informed when an Address Resolution Protocol (ARP) check was performed successfully. As a consequence, the users could be confused about the time needed to load the interface. With this update, a message is printed after every ARP check thus preventing confusion.

BZ#921476

Previously, initscripts documentation contained no information about the rule-* files. As a consequence, users did not know how to set routing rules for IPv6 addresses. This update adds documentation for the rule6-* files to the sysconfig.txt file.

BZ#905423

Previously, users were not aware of the /etc/init/*.conf files being overwritten after every update with the default values. A comment has been added to the /etc/init/*.conf files to inform users these files should not be modified and to use the *.override files instead.

Enhancements**BZ#815676**

With this update, configuration options for Dynamic Host Configuration Protocol version 6 (DHCPv6) have been applied to the /etc/dhcp/dhclient6-<iface>.conf files. Options for both DHCPv4 and DHCPv6 in the /etc/dhcp/dhclient6-<iface>.conf are now applied.

Users of initscripts are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.70. IOTOP

8.70.1. RHBA-2013:1719 — iotop bug fix update

An updated iotop package that fixes several bugs is now available for Red Hat Enterprise Linux 6.

The iotop package provides a program with a UI similar to the "top" utility. The program watches input-output (I/O) usage information output by the Linux kernel and displays a table of current I/O usage by processes on the system.

Bug Fixes

BZ#746240, BZ#908149

Previously, the iotop utility terminated unexpectedly when it was run by a non-root user. This was because a recently-applied patch in CVE-2011-2494 made I/O statistics from the taskstats kernel subsystem accessible only to root users, and iotop did not anticipate that its "taskstats" call could fail when run by a non-root user. This update adds permission checks to iotop, and when the user does not have the necessary permissions, iotop exits with an explanation that root privileges are now required.

BZ#826875

Previously, the iotop utility did not handle platform strings correctly. Consequently, the iotop command could not show the I/O scheduling class and its priority ("PRIO") column on 64-bit PowerPC systems properly. With this update, the bug has been fixed so that the iotop command now shows the "PRIO" column on 64-bit PowerPC systems as expected.

BZ#849559

When an invalid locale was set, the iotop utility failed to start with the following traceback error:

```
locale.Error: unsupported locale setting
```

With this update, the underlying source code has been modified. As a result, when an invalid locale is set, the default locale is used instead and a warning about this change is returned.

Users of iotop are advised to upgrade to this updated package, which fixes these bugs.

8.71. IPA

8.71.1. RHBA-2013:1651 — ipa bug fix and enhancement update

Updated ipa packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

Red Hat Identity Management is a centralized authentication, identity management and authorization solution for both traditional and cloud-based enterprise environments. It integrates components of the Red Hat Directory Server, MIT Kerberos, Red Hat Certificate System, NTP, and DNS. It provides web browser and command-line interfaces. Its administration tools allow an administrator to quickly install, set up, and administer a group of domain controllers to meet the authentication and identity management requirements of large-scale Linux and UNIX deployments.

Bug Fixes

BZ#904119

Previously, during migration, users were added to the default user group one by one. As a consequence, adding users to a large group was time consuming. With this update, users are now added in batches of 100, which provides a considerable performance boost over the previous method.

BZ#905626

Previously, the Identity Management client installer did not look for all available servers when it tried to enroll a client. Consequently, the enrollment "ipa-client-install" command failed to enroll a client if any of the Identity Management masters were unavailable during the enrollment. With this update, the client installer tries all servers, either auto-discovered from DNS or passed using the "--server" option on the command line, until it finds an available server, and ipa-client-install now works properly.

BZ#906846

Identity Management did not work correctly when migrating from an OpenLDAP server. As a consequence, attempts to retrieve the LDAP schema from the remote server failed. With this update, Identity Management also looks in the "cn=subschema" entry, and migrations from OpenLDAP servers no longer fail.

BZ#907881

Prior to this update, the Identity Management password lockout Directory Server plug-in processed password lockout incorrectly. Consequently, if Identity Management password policy was configured with the Lockout Time value set to 0, user accounts were permanently disabled even though the maximum number of user password failures had not been exceeded. The plug-in has been fixed to process the password lockout time correctly, and the user accounts lockout now works as expected.

BZ#915745

Previously, update files used when upgrading an Identity Management server to a later version did not contain the new Directory Server schema "ipaExternalMember" attribute type and the "ipaExternalGroup" object class. Consequently, neither command-line interface (CLI) commands using the schema elements nor web user interface (Web UI) as a whole worked correctly. This update adds the missing object class and attribute type to the Identity Management update files. The Directory Server schema is now updated during the Identity Management update process, and both CLI commands and the Web UI work properly.

BZ#916209

The Identity Management configuration parser was not able to parse the Kerberos client configuration file (/etc/krb5.conf) when it contained the "includedir" directive. Consequently, the Identity Management ipa-adtrust-install installer, which directly parses and updates Kerberos client configuration, terminated unexpectedly with a syntax error. With this update, the configuration parser processes "includedir" correctly, and ipa-adtrust-install no longer crashes in the described scenario.

BZ#924004

Previously, when the Identity Management client installer was downloading a Certification Authority (CA) certificate from Identity Management server using the LDAP protocol, it did occasionally not fallback to the HTTP protocol. Consequently, Identity Management client installation failed even though the certificate was accessible using the HTTP protocol. With this update, the Identity Management client installer can properly fallback between different protocols when downloading a CA certificate, and it is now able to complete the installation even when download via one protocol fails.

BZ#924009

The Identity Management client installer did not allow re-enrolling of an already enrolled client. Consequently, when a machine or a virtual machine with a configured Identity Management client was being removed or decommissioned without unenrolling the client first, all succeeding client enrollments failed until the client entry was removed from the Identity Management sever. This update adds a "--force-join" option to the Identity Management client installer, and the privileged administrator is now able to re-enroll an Identity Management client.

BZ#924542

Previously, Identity Management Host Based Access Control (HBAC) rules API allowed administrators to specify a "Source Host" component of HBAC rules even though this component had been deprecated. Consequently, unexpected behavior could occur when using the "Source Host" component in HBAC rules. This bug has been fixed; "Source Host" components are now not allowed in HBAC rules, and unexpected behavior of the rules for administrators no longer occurs.

BZ#948928

Under certain circumstances, the Identity Management upgrade process double encoded the Certification Authority (CA) certificate stored in Directory Server. Consequently, some Identity Management clients failed to decode the CA certificate and installing a client failed. With this update, CA certificates are now properly encoded; client installation CA certificate is correctly retrieved from Identity Management server and the installation proceeds as expected.

BZ#950014

In some cases, the Identity Management installation and upgrade process did not update the user and user role membership information in correct order. As a consequence, user roles were occasionally not correctly applied, and users could fail to proceed with privileged actions even though they had been authorized for them (for example, enrollment of an an Identity Management client). Now, the membership information is applied in correct order, and users' privileged actions no longer fail because of incomplete membership information.

BZ#952241

Previously, when an Identity Management public-key infrastructure (PKI) server certificate (auditSigningCert) was being renewed, incorrect trust argument was assigned to the renewed certificate and the server was unable to use it. The certificate renewal procedure has been updated to assign correct trust arguments to the renewed certificates, and Identity Management PKI certificate renewal now works as expected.

BZ#967870

Identity Management server with Active Directory integration support configured replies differently in NetLogon queries compared to Active Directory. The following discrepancies were present in NetLogon behavior:

- No response to NetLogon query when querying over TCP based LDAP
- No response when DnsDomain was not present in the query
- No return of a LDAP_RES_SEARCH_RESULT to sender when query did not match; NetLogon became unresponsive.

As a consequence, these discrepancies could cause errors in utilities which had sent the NetLogon queries. The NetLogon query responder has been fixed, and the above mentioned issues in NetLogon replies no longer occur.

BZ#970541

Identity Management server did not work efficiently in case of entries with many members, such as a large user group. Consequently, Identity Management CLI or Web UI management commands operating with such entries (for example, adding new users, listing groups, or updating them) could last more than 30 seconds. Several improvements have been implemented in the Identity Management server, namely:

- Web UI interface now avoids membership information when it is not required (for example, in group listing)
- Entry membership manipulating commands (for example, adding users to a group) now avoid unnecessary manipulation with membership information
- Missing substring indices for membership attributes have been added.

With these implementations, the performance of Identity Management CLI and Web UI management commands has been significantly improved, especially when dealing with large user groups.

BZ#975431

Previously, the `/var/lib/ipa/pki-ca/publish/` directory, where Identity Management public-key infrastructure (PKI) publishes Certificate Revocation List (CRL) exports, contained incorrect ownership and permissions information after the `ipa-server` package had been reinstalled or upgraded. Consequently, PKI was not able to update CRL in the directory until the ownership and permissions of the directory were manually amended. The Identity Management installer and upgrade script have been fixed to handle the ownership and permissions of the directory correctly, and CRL exports are now updated properly in the described scenario.

BZ#976716

Prior to this update, the Identity Management XML-RPC interface occasionally did not return the correct "Content-Type" header in its replies. Consequently, programs or scripts processing the XML-RPC response could fail to process the response with a validation error. The XML-RPC responder has been fixed to return the correct "Content-Type" header, and programs and scripts are now able to call the Identity Management XML-RPC interface even with strict validation enabled.

BZ#980409

Previously, the Identity Management Active Directory integration did not expect different procedure for populating `KERB_VALIDATION_INFO` section of MS-PAC extension for a Kerberos ticket done in Microsoft Windows Server 2012, as compared to Microsoft Windows Server 2008. As a consequence, such Kerberos tickets were not accepted due to an incompatibility and could not be used to authenticate or create a Trust with the Microsoft Windows Server 2012. The `KERB_VALIDATION_INFO` verification has been refactored to filter out unexpected values before further processing, and the Identity Management Active Directory Trust creation no longer fails with Microsoft Windows Server 2012.

BZ#1011044

Previously, the `ipa-client-install` installation script did not properly detect whether the client had already been installed on the machine or not. As a consequence, the client uninstall script could refuse to restore the machine when it did not recognize the client as installed. Also, the client installation could succeed even on an installed Identity Management client or server machine. This, however, could disrupt the configuration files or Identity management client or server function. With this update, `ipa-client-install` has been fixed to detect installation properly, and the issues described above no longer occur.

Enhancements

BZ#955698

This update introduces the "userClass" attribute for Identity Management server host entries. Previously, host entries did not contain a free-form attribute usable for host provisioning systems to tag or set a class for a new host, which could then be used by other functions of Identity Management, for example, by the Automatic Membership Assignment module. Administrators and provisioning systems are now able to use the new "userClass" host entry attribute.

BZ#986211

This update adds the "GECOS" field for user entries to Identity Management Web UI. "GECOS" is an important user field as it equals the user's common name presented to the systems, and it should be editable both through CLI and Web UI interfaces. Now, the user's "GECOS" field can be displayed and changed in Identity Management Web UI.

Users of ipa are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.72. IPMITOOL

8.72.1. RHBA-2013:1259 — ipmitool bug fix update

Updated ipmitool packages that fix several bugs are now available.

The ipmitool package contains a command-line utility for interfacing with devices that support the Intelligent Platform Management Interface (IPMI) specification. IPMI is an open standard for machine health, inventory, and remote power control.

Bug Fixes

BZ#826027

In a previous ipmitool update, the new options "-R" and "-N" were added to adjust the retransmission rate of outgoing IPMI requests over LAN and lanplus interfaces. Implementation of these options set a wrong default value of the retransmission timeout, and an outgoing request timed out prematurely. In addition, in some corner cases, ipmitool could terminate unexpectedly with a segmentation fault when the timeout occurred. This update fixes the default timeout value, so ipmitool without the "-N" option retransmits outgoing IPMI requests as in previous versions, and crashes no longer occur in the described scenario.

BZ#903251

Previously, enabling the "ipmi" and "link" keys in user access information using the ipmitool utility did not work properly. Consequently, the values of these settings were not taken into account. A patch has been provided that ensures the values of these settings are read and processed as expected.

BZ#923192

In cases of congested network or slow-responding Baseboard Management Controller (BMC), the reply operation timeout triggered the protocol command retry action. Consequently, the ipmitool utility could incorrectly process a LAN session protocol command with the reply from a previous protocol command. This update fixes handling of expected replies for each command alone, and cleans up expected replies between commands. Now, the retried reply of the first command is correctly ignored while the later command, which is currently pending, is properly processed in the described scenario.

Users of `ipmitool` are advised to upgrade to these updated packages, which fix these bugs. After installing this update, the IPMI event daemon (`ipmievtd`) will be restarted automatically.

8.73. IPROUTE

8.73.1. RHBA-2013:1697 — iproute bug fix and enhancement update

Updated `iproute` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `iproute` packages contain networking utilities (for example, `ip` and `rtmon`), which are designed to use the advanced networking capabilities of the Linux kernel.

Bug Fixes

BZ#1011148

While monitoring IP neighbor cache with the `ip monitor neigh` command, the cache experienced the layer 2 network miss. Consequently, `ip monitor neigh` command could not decode the miss event generated by the kernel. To fix this bug, code for neighbor cache events for entry deletion and entry miss have been back-ported from upstream and `ip monitor neigh` now recognizes cache miss event and format it properly with a `miss` keyword on the output.

BZ#950400

Previously, Red Hat Enterprise Linux 6 was missing a functionality to set up IPv6 token-only network configuration. As a consequence, the user had fewer networking options. The IPv6 token feature has been implemented in both kernel (BZ#876634) and a userspace interface to `iproute`. Users can now setup IPv6 token-only networking, optionally receiving network prefixes later.

BZ#908155

Red Hat Enterprise Linux 6.5 shipped with VXLAN (Virtual Extended LAN), a VLAN-like layer 3 encapsulation technique support in the kernel, so a userspace interface was required for users and applications to utilize the VXLAN feature. With this update, the `ip` utility recognizes and supports the 'vxlan' devices.

BZ#838482

When larger `rto_min` (the minimum TCP Retransmission TimeOut to use when communicating with a certain destination) was set, the `ip route show` command did not return correct values. A patch has been provided to fix this bug and `ip route show` now handles `rto_min` as expected.

BZ#974694

Prior to this update, the manual page for the `lstat` utility was referring wrongly to non-existent directory, the `iproute-doc` instead of `iproute-<package version>` directory. The incorrect documentation could confuse the user. To fix this bug, the file-system path has been corrected.

BZ#977845

Previously, there was an inconsistency between the `lstat` utility's interval option behavior and its documentation. Consequently, `lstat` exited after a number of seconds instead of refreshing the view, making the interval option useless. The interval option behavior has been changed to refresh the data every `N` seconds, thus fixing the bug.

BZ#985526

Previously, the **ip** utility was mishandling netlink communication, which could cause hangs under certain circumstances. Consequently, listing network devices with the **ip link show** command hung in a SELinux restricted mode. With this update, the **ip** utility checks for the result of the **rtnl_send()** function before waiting for a reply, avoiding an indefinite hang. As a result, it is now possible to list network devices in a SELinux restricted environment.

BZ#950122

Prior to this update, the **tc** utility documentation lacked description of the **batch** option. To fix this bug, the **tc** manual pages have been updated including the description of the **batch** option.

Enhancements**BZ#885977**

Previously, the bridge module **sysfs** system did not provide the ability to inspect the non-configuration IP multicast Internet Group Management Protocol (IGMP) snooping data. Without this functionality, users could not fully analyze their multicast traffic. With this update, users are able to list detected multicast router ports, groups with active subscribers and the associated interfaces.

BZ#929313

Distributed Overlay Virtual Ethernet (DOVE) tunnels allow for building of Virtual Extensible Local Area Network (VXLAN), which represents a scalable solution for ISO OSI layer 2 networks used in cloud centers. The bridge tool is part of the **iproute** packages and can be used, for example, to manage forwarding database on WLAN devices on Linux platform.

BZ#851371

If the **tc** utility is instrumented from a pipe, there is no way how to recognize when a subcommand has been completed. A new **OK** option has been added to the **tc** utility. Now, **tc** in the batch mode accepts commands in standard input (the **tc -OK -force -batch** command) and returns **OK** on a new line on standard output for each successfully completed **tc** subcommand.

Users of **iproute** are advised to upgrade to these updated packages, which fixe these bugs and add these enhancements.

8.74. IPTABLES

8.74.1. RHBA-2013:1710 — iptables bug fix and enhancement update

Updated iptables packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The iptables utility controls the network packet filtering code in the Linux kernel. The utility allows users to perform certain operations such as setting up firewalls or IP masquerading.

Bug Fixes**BZ#924362**

A previous version of iptables added the "alternatives" functionality support for the **/lib/xtables/** or **/lib64/xtables/** directory. However, iptables failed to replace the directory with the alternatives slave

symbolic link when upgrading iptables with the "yum upgrade" command and the directory contained custom plug-in files. Consequently, some iptables modules became unavailable. This problem has been fixed by modifying the iptables spec file so that the /lib/xtables/ or /lib64/xtables/ directory is no longer managed by "alternatives".

BZ#983198

The iptables-save command previously supported only the "--modprobe=" option to specify the path to the modprobe executable. However, the iptables-save(8) man page incorrectly stated that this action could have been performed using an unsupported option, "-M", which could lead to confusion. The iptables-save command has been modified to support the "-M" option for specifying the path to modprobe, and corrects the iptables-save(8) man page, which now correctly mentions both the "-M" and "--modprobe=" option.

BZ#1007632

Due to a bug in the iptables init script, the system could become unresponsive during shutdown when using the network-based root device and the default filter for INPUT or OUTPUT policy was DROP. This problem has been fixed by setting the default chain policy to ACCEPT before flushing the iptables rules and deleting the iptables chains.

Enhancements**BZ#845435**

The iptables utility has been modified to support a new option, "--queue-bypass", which allows bypassing an NFQUEUE rule if the specified queue is not used.

BZ#928812

A new iptables service option, "reload", has been added to enable a refresh of the firewall rules without unloading netfilter kernel modules and a possible drop of connections.

Users of iptables are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.75. IPVSDM

8.75.1. RHBA-2013:1639 — ipvsadm bug fix and enhancement update

Updated ipvsadm packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The ipvsadm package provides the ipvsadm tool to administer the IP Virtual Server services offered by the Linux kernel.

**NOTE**

The ipvsadm packages have been upgraded to upstream version 1.26, which provides new features to IPVS from the kernel side, which take full-advantage of PE config of SIP PE-data. In addition, this update:

- * fixes the "One Packet Scheduler" output for the status and save operations to include all Virtual Servers instead of only those configured with a persistent flag (BZ#986189)
- * addresses a possible, but very unlikely, memory corruption issues;
- * includes minor improvements to the manual pages related to ipvsadm.

Users of ipvsadm are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.76. IRQBALANCE

8.76.1. RHBA-2013:1583 — irqbalance bug fix update

Updated irqbalance packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The irqbalance packages provide a daemon that evenly distributes interrupt request (IRQ) load across multiple CPUs for enhanced performance.

Bug Fixes

BZ#951720

Previously, irqbalance warned about MSI interrupts, and that IRQs would not be properly classified due to the use of a kernel version older than kernel-2.6.32-279. This update blocks users from using irqbalance with an older version of the kernel, without features required for processing MSI interrupts, and warning messages are no longer received.

BZ#975524

Due to recent changes in the irqbalance packages, the `/var/run/irqbalance.pid` file was not created upon start of the irqbalance service, causing irqbalance to become non-compliant with the Linux Standard Base (LSB) specification. This update provides a patch fixing this problem so the irqbalance packages are LSB compliant again.

BZ#991363

A bug in the irqbalance code caused the irqbalance daemon to terminate with a segmentation fault when a CPU was hot plugged or hot unplugged. This update fixes a corrupted IRQ rebalance list and the irqbalance daemon no longer crashes in this scenario.

Users of irqbalance are advised to upgrade to these updated packages, which fix these bugs.

8.77. ISCSI-INITIATOR-UTILS

8.77.1. RHBA-2013:1700 — iscsi-initiator-utils bug fix and enhancement update

Updated iscsi-initiator-utils packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The iscsi-initiator-utils packages provide the server daemon for the iSCSI protocol, as well as utilities used to manage the daemon. iSCSI (Internet Small Computer System Interface) is a protocol for distributed disk access using SCSI commands sent over Internet Protocol networks.



NOTE

The iscsi-initiator-utils packages have been upgraded to upstream version 6.2.0.873, which provides a number of bug fixes and enhancements over the previous version. (BZ#[916007](#))

Bug Fixes

BZ#[884427](#)

Previously, database errors could occur if multiple node records in different formats were created for the same iSCSI target portal. Consequently, depending on the file system dependent return order of the readdir syscall, an error occasionally occurred causing an update operation to fail. To fix this bug, multiple node records in different formats have been prevented from existing simultaneously and detected at record creation time. Duplicate node entries no longer exist in the iSCSI database, and updates to records do not result in database errors.

BZ#[983553](#)

Prior to this update, a single unreachable target could previously block rescans of others. Consequently, the iscsiadm utility could halt in the D state and the rest of the targets could remain unscanned. To fix this bug, iscsiadm has been made terminable and all the targets have been updated. Now, functioning sessions will be rescanned properly without long delays.

BZ#[1001705](#)

When VDMS (Virtual Desktop Server Manager) attempted to add a new record to the iSCSI database, it failed with the following error:

```
iscsiadm: Error while adding record: no available memory.
```

Consequently, due to this error, the host became non-operational when connecting to storage. An upstream patch has been applied and the /var/lib/iscsi file is now successfully attached.

Enhancements

BZ#[831003](#)

For the bnx2i hardware and potentially other offloading solutions (complementary network technologies for delivering data originally targeted for cellular networks), the iscsistart tool for passing along the VLAN tag from iBFT (iSCSI Boot Firmware Table) to iface_rec (iscsi iface record name) has been implemented to this package.

BZ#[917600](#)

With this update, support for managing Flash nodes from the open-iscsi utility has been added to this package.

Users of `iscsi-initiator-utils` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.78. IW

8.78.1. RHEA-2013:1563 — iw enhancement update

Updated `iw` packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

The `iw` command-line utility is used for configuring wireless devices based on the `nl80211` interface.



NOTE

The `iw` packages have been upgraded to upstream version 3.10, which provides one enhancement over the previous version. This update adds support for Wake on Wireless LAN (WoWLAN) to Atheros WiFi interfaces in Red Hat Enterprise Linux 6. (BZ#[951706](#))

Users of `iw` are advised to upgrade to these updated packages, which add this enhancement.

8.79. JAVA-1.6.0-OPENJDK

8.79.1. RHBA-2013:1741 — java-1.6.0-openjdk bug fix and enhancement update

Updated `java-1.6.0-openjdk` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `java-1.6.0-openjdk` packages provide the OpenJDK 6 Java Runtime Environment and the OpenJDK 6 Java Software Development Kit.



NOTE

The `java-1.6.0-openjdk` packages have been upgraded to upstream IcedTea version 1.13.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#[983411](#))

Bug Fix

BZ#[976897](#)

Previously, `int[]` objects allocated by instances of the `com.sun.imageio.plugins.jpeg.JPEGImageWriter` class were consuming extensive amounts of memory, which was consequently not released. With this update, the underlying stream processing logic has been modified to ensure correct releasing of such memory, and extensive memory consumption no longer occurs.

Users of `java-1.6.0-openjdk` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. All running instances of OpenJDK Java must be restarted for the update to take effect.

8.80. JAVA-1.7.0-OPENJDK

8.80.1. RHBA-2013:1611 — java-1.7.0-openjdk bug fix and enhancement update

Updated java-1.7.0-openjdk packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The java-1.7.0-openjdk packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Software Development Kit.



NOTE

The java-1.7.0-openjdk package has been upgraded to upstream version 2.4.3, which provides a number of bug fixes and enhancements over the previous version.

Bug Fixes

BZ#825824

Attempting to compile a SystemTap script using the jstack tapset could have failed with an error similar to the following:

```
error: the frame size of 272 bytes is larger than 256 bytes
```

This update corrects the jstack tapset and resolves this problem.

BZ#871771

Because of incorrect KDC list concatenation logic, the `sun.security.krb5.Config.getKDCList` method returned incorrect KDC lists when the `dns_lookup_kdc` property in the `krb5.conf` file was set to `true`. The concatenation logic has been fixed with this release and correct KDC lists are now returned.

BZ#997633

The java-1.7.0-openjdk RPM package contained incorrect specification of the `libnss3` dependency and installed its `x86_64` version on `i686` systems. Because of the missing dependency, launching the `java` command with the `-Dcom.sun.management.jmxremote` parameter on 32bit JVMs terminated unexpectedly. The dependency specification has been corrected with this update. As a result, the correct version of the `libnss3` package is installed and the `java` command no longer terminates when launched with the `-Dcom.sun.management.jmxremote` parameter.

Enhancements

BZ#831734, BZ#905128

The NSS security provider is now the default security provider in OpenJDK 7. This brings a significant performance improvement over the previous releases.

BZ#916288

The java-1.7.0-openjdk RPM package now provides the `java` dependency. As a result, it is no longer necessary to have the java-1.6.0-openjdk package installed alongside java-1.7.0-openjdk for the `java` dependency to be available.

Users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. All running instances of OpenJDK Java must be restarted for the update to take effect.

8.81. KDE-SETTINGS

8.81.1. [RHBA-2013:1053](#) — [kde-settings bug fix update](#)

Updated kde-settings packages that fix one bug are now available.

The kde-settings packages provide a rich set of administration panels to configure system and desktop settings in the Konqueror Desktop Environment (KDE).

Bug Fix

BZ#886237

The Konqueror browser enabled Java support by default. Because Java is one of the common targets for browser-based malware attacks, Java is now disabled by default in Konqueror.

To enable Java in Konqueror, navigate to Settings -> Configure Konqueror -> Java & JavaScript (which sets the path to Java), and select the "Enable Java globally" check box.

Users of kde-settings are advised to upgrade to these updated packages, which fix this bug.

8.82. KERNEL

8.82.1. [RHSA-2015:0062](#) — [Important: kernel security, bug fix, and enhancement update](#)

Updated kernel packages that fix multiple security issues, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.5 Extended Update Support.

Red Hat Product Security has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links in the References section.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

[CVE-2014-3673](#), [CVE-2014-3687](#), Important

A flaw was found in the way the Linux kernel's SCTP implementation handled malformed or duplicate Address Configuration Change Chunks (ASCONF). A remote attacker could use either of these flaws to crash the system.

[CVE-2014-3688](#), Important

A flaw was found in the way the Linux kernel's SCTP implementation handled the association's output queue. A remote attacker could send specially crafted packets that would cause the system to use an excessive amount of memory, leading to a denial of service.

[CVE-2014-5045](#), Moderate

A flaw was found in the way the Linux kernel's VFS subsystem handled reference counting when performing unmount operations on symbolic links. A local, unprivileged user could use this flaw to exhaust all available memory on the system or, potentially, trigger a use-after-free error, resulting in a system crash or privilege escalation.

[CVE-2014-4608](#), Low

An integer overflow flaw was found in the way the `lzo1x_decompress_safe()` function of the Linux kernel's LZO implementation processed Literal Runs. A local attacker could, in extremely rare cases, use this flaw to crash the system or, potentially, escalate their privileges on the system.

Red Hat would like to thank Vasily Averin of Parallels for reporting CVE-2014-5045, and Don A. Bailey from Lab Mouse Security for reporting CVE-2014-4608. The CVE-2014-3673 issue was discovered by Liu Wei of Red Hat.

Bug Fixes

BZ#1108360

Before this update, under certain conditions, the kernel timer could cause the Intelligent Platform Management Interface (IPMI) driver to become unresponsive, resulting in high CPU load. With this update, a patch has been applied, and the IPMI driver no longer hangs.

BZ#1109270, BZ#1109712

Previously, when error recovery was restarted, the Orthonormal Basis Functions (OBF) timer in the KCS driver was not reset, which led to an immediate timeout. As a consequence, these timing issues caused `ipmi` to become unresponsive. In addition, numerous error messages were filling up the `/var/log/messages` file and causing high CPU usage. With this update, patches have been applied to fix this bug, and `ipmi` no longer hangs in the described situation.

BZ#1135993

Due to certain kernel changes, the TCP Small Queues (TSQ) process did not handle Nagle's algorithm properly when a TCP session became throttled. The underlying source code has been patched, and Nagle's algorithm now works correctly in TSQ.

BZ#1140976

Before this update, due to a bug in the error-handling path, corrupted metadata block could be used as a valid block. With this update, the error handling path is fixed and more checks are added to verify the metadata block. Now, when a corrupted metadata block is encountered, it is properly marked as corrupted and handled accordingly.

BZ#1154087, BZ#1158321

Previously, log forces with relatively little free stack available occurred deep in the call chain. As a consequence, a stack overflow in the (XFS) file system and the system could terminate unexpectedly. To fix this bug, moving log forces to a work queue relieves the stack pressure and avoids the system crash.

BZ#1158324

Before this update, TCP transmit interrupts could not be set lower than the default of 8 buffered tx frames, which under certain conditions led to TCP transmit delays occurring on ixgbe adapters. With this update, code change removes the restriction of minimum 8 buffered frames and now allows minimum of 1 frame a transmit to occur. And as a result, transmit delays are now minimized.

BZ#1165984

Previously, a coding error in Ethernet 100 driver update caused improper initialization for certain Physical Layers (PHYs) and return of RX errors. With this update, the coding error has been fixed, and the device driver works properly.

BZ#1158327

Before this update, the frame buffer (offb) driver did not support setting of the color palette registers on the QEMU standard VGA adapter, which caused incorrect color displaying. The offb driver has been updated for the QEMU standard VGA adapter, fixing the color issues.

BZ#1142569

Before this update, several race conditions occurred between PCI error recovery callbacks and potential calls of the ifup and ifdown commands in the tg3 driver. When triggered, these race conditions could cause unexpected kernel termination. This bug has been fixed, and the kernel no longer crashes.

BZ#1158889, BZ#1162748

Due to hardware bug conditions during Top Segmentation Offload (TSO) fragment processing, there was a page allocation failure in kernel and packets were not transmitted. With this update, more generic Generic Segmentation Offload (GSO) is used as a fallback when TSO fragment processing fails, and packets are now successfully transmitted.

BZ#1163397

Previously, the kernel became unresponsive when using a zombie PID and cgroup. To fix this bug, a patch has been applied, and the kernel no longer hangs.

BZ#1165000

Previously, under certain error conditions gfs2_converter introduced incorrect values for the on-disk inode's di_goal_meta field. As a consequence, gfs2_converter returned the EBADSLT error on such inodes and did not allow creation of the new files in directories or new blocks in regular files. The fix allows gfs2_converter to set a sensible goal value if a corrupt one is encountered and proceed with normal operations. With this update, gfs2_converter implicitly fixes any corrupt goal values, and thus no longer disrupts normal operations.

BZ#1169403

Previously, certain error conditions led to messages being sent to system logs. These messages could become lost instead of being logged, or repeated messages were not suppressed. In extreme cases, the resulting logging volume could cause system lockups or other problems. The relevant test has been reversed to fix this bug, and frequent messages are now suppressed and infrequent messages logged as expected.

Enhancement**BZ#1167209**

This update adds fixes from Emulex and Oracle Enterprise Management (OEM) qualifications including latest fixes for Skyhawk hardware to the Emulex be2iscsi driver.

Users of kernel are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add this enhancement. The system must be rebooted for this update to take effect.

8.82.2. RHSA-2014:1668 — Important: kernel security, bug fix, and enhancement update

Updated kernel packages that fix one security issue, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.5 Extended Update Support.

The Red Hat Security Response Team has rated this update as having Important security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available from the CVE link associated with the description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2014-5077, Important

A NULL pointer dereference flaw was found in the way the Linux kernel's Stream Control Transmission Protocol (SCTP) implementation handled simultaneous connections between the same hosts. A remote attacker could use this flaw to crash the system.

Bug Fixes

BZ#1110839

Due to a bug in the kernel signal handling, the decimal floating point (DFP) operations could have been executed with an incorrect rounding mode. As a consequence, DFP calculations could return incorrect or corrupted results. This update fixes this problem by replacing a simple bit mask that was previously used to verify validity of some values in the floating point control register. The bit mask is replaced by a trial load of the floating point control register.

BZ#1140163

Previously, when freeing a large number of huge pages (several TB), the kernel could experience soft lockup events. This could possibly result in performance problems. The memory management code has been modified to increase a chance of a context switch in this situation, which prevents occurrence of soft lockup events.

BZ#1122102

A bug in the nouveau driver could prevent the main display of a Lenovo ThinkPad W530 laptop from being initialized after the system was resumed from suspend. This happened if the laptop had an external screen that was detached while the system was suspended. This problem has been fixed by backporting an upstream patch related to the DisplayPort interface.

BZ#1139807

Due to race conditions in the IP Virtual server (IPVS) code, the kernel could trigger a general protection fault when running the IPVS connection synchronization daemon. With this update, the race conditions in the IPVS code have been addressed, and the kernel no longer crashes when running the IPVS daemon.

BZ#1139345

The kernel could sometimes panic due to a possible division by zero in the kernel scheduler. This bug has been fixed by defining a new `div64_ul()` division function and correcting the affected calculation in the `proc_sched_show_task()` function.

BZ#1125980

Removing the `rtsc_pci_ms` kernel module on some Lenovo ThinkPad series laptops could result in a kernel panic. This update resolves this problem by correcting a bug in the base drivers function, `platform_uevent()`.

BZ#1125994

A bug in the Linux Netpoll API could result in a kernel oops if the system had the netconsole service configured over a bonding device. With this update, incorrect flag usage in the netpoll_poll_dev() function has been fixed and the kernel no longer crashes due to this bug.

BZ#1127580

The kernel did not handle exceptions caused by an invalid floating point control (FPC) register, resulting in a kernel oops. This problem has been fixed by placing the label to handle these exceptions to the correct place in the code.

BZ#1138301

Previously, certain network device drivers did not accept ethtool commands right after they were mounted. As a consequence, the current setting of the specified device driver was not applied and an error message was returned. The ETHTOOL_DELAY variable has been added, which makes sure the ethtool utility waits for some time before it tries to apply the options settings, thus fixing the bug.

BZ#1130630

A rare race between the file system unmount code and the file system notification code could lead to a kernel panic. With this update, a series of patches has been applied to the kernel to prevent this problem.

BZ#1131137

A bug in the bio layer could prevent user space programs from writing data to disk when the system run under heavy RAM memory fragmentation conditions. This problem has been fixed by modifying a respective function in the bio layer to refuse to add a new memory page only if the page would start a new memory segment and the maximum number of memory segments has already been reached.

BZ#1135713

Due to a bug in the ext3 code, the fdatsync() system call did not force the inode size change to be written to the disk if it was the only metadata change in the file. This could result in the wrong inode size and possible data loss if the system terminated unexpectedly. The code handling inode updates has been fixed and fdatsync() now writes data to the disk as expected in this situation.

BZ#1134258

Previously, the openvswitch driver did not handle frames that contained multiple VLAN headers correctly, which could result in a kernel panic. This update fixes the problem and ensures that openvswitch process such frames correctly.

BZ#1134696

Later Intel CPUs added a new "Condition Changed" bit to the MSR_CORE_PERF_GLOBAL_STATUS register. Previously, the kernel falsely assumed that this bit indicates a performance interrupt, which prevented other NMI handlers from running and executing. To fix this problem, a patch has been applied to the kernel to ignore this bit in the perf code, enabling other NMI handlers to run.

BZ#1135393

After the VLAN devices over the virtio_net driver were allowed to use the TCP Segmentation Offload (TSO) feature, the segmentation of packets was moved from virtual machines to the host. However, some devices cannot handle TSO using the 8021q module, and are breaking the packets, which resulted in very low throughput (less than 1 Mbps) and transmission of broken packets over the wire.

Until this problem is properly fixed, a patch that allows using of the TSO feature has been reverted; the segmentation is now performed again on virtual machines as and the network throughput is normal.

BZ#1141165

Due to a race condition in the IP Virtual server (IPVS) code, the kernel could trigger a panic when processing packets from the same connection on different CPUs. This update adds missing spin locks to the code that hashes and unhashes connections from the connection table, and ensures that all packets from the same connection are processed by a single CPU.

BZ#1129994

Previously, small block random I/O operations on IBM Power 8 machines using Emulex 16 Gb Fibre Channel (FC) Host Bus Adapter (HBA) could become unresponsive due to a bug in the lpfc driver. To fix this problem, a memory barrier has been added to the lpfc code to ensure that a valid bit is read before the CQE payload.

BZ#1126681

Running the "bridge link show" command on a system with configured bridge devices could trigger a kernel panic. This happened because all RTNL message types were not properly unregistered from the bridge module registers. This update ensures that both RTNL message types are correctly unregistered and the kernel panic no longer occurs in this situation.

BZ#1114406

Previously, the NFS server did not handle correctly situations when multiple NFS clients were appending data to a file using write delegations, and the data might become corrupted. This update fixes this bug by adjusting a NFS cache validity check in the relevant NFS code, and the file accessed in this scenario now contains valid data.

BZ#1131977

Previously, the IPv4 routing code allowed the IPv4 garbage collector to run in parallel on multiple CPUs with the exact configuration. This could greatly decrease performance of the system, and eventually result in soft lockups after the system reached certain load. To resolve this problem and improve performance of the garbage collector, the collector has been moved to the work queue where it is run asynchronously.

Enhancements**BZ#1133834**

A new "nordirplus" option has been implemented for the exportfs utility for NFSv3. This option allows the user to disable REaddirPLUS requests for the given NFSv3 export, and thus prevent unwanted disk access in certain scenarios.

All kernel users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add this enhancement. The system must be rebooted for this update to take effect.

8.82.3. RHSA-2014:1167 — Important: kernel security and bug fix update

Updated kernel packages that fix multiple security issues, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2014-0205, Important

A flaw was found in the way the Linux kernel's futex subsystem handled reference counting when requeuing futexes during `futex_wait()`. A local, unprivileged user could use this flaw to zero out the reference counter of an inode or an mm struct that backs up the memory area of the futex, which could lead to a use-after-free flaw, resulting in a system crash or, potentially, privilege escalation.

CVE-2014-3535, Important

A NULL pointer dereference flaw was found in the way the Linux kernel's networking implementation handled logging while processing certain invalid packets coming in via a VxLAN interface. A remote attacker could use this flaw to crash the system by sending a specially crafted packet to such an interface.

CVE-2014-3917, Moderate

An out-of-bounds memory access flaw was found in the Linux kernel's system call auditing implementation. On a system with existing audit rules defined, a local, unprivileged user could use this flaw to leak kernel memory to user space or, potentially, crash the system.

CVE-2014-4667, Moderate

An integer underflow flaw was found in the way the Linux kernel's Stream Control Transmission Protocol (SCTP) implementation processed certain `COOKIE_ECHO` packets. By sending a specially crafted SCTP packet, a remote attacker could use this flaw to prevent legitimate connections to a particular SCTP server socket to be made.

Red Hat would like to thank Gopal Reddy Kodudula of Nokia Siemens Networks for reporting CVE-2014-4667. The security impact of the CVE-2014-0205 issue was discovered by Mateusz Guzik of Red Hat.

Bug Fixes

BZ#1089359

Previously, NFSv4 allowed an NFSv4 client to resume an expired or lost file lock. This could result in file corruption if the file was modified in the meantime. This problem has been resolved by a series of patches ensuring that an NFSv4 client no longer attempts to recover expired or lost file locks.

BZ#1090613

A false positive bug in the NFSv4 code could result in a situation where an `NFS4ERR_BAD_STATEID` error was being resent in an infinite loop instead of a bad state ID being recovered. To fix this problem, a series of patches has been applied to the NFSv4 code. The NFS client no longer retries an I/O operation that resulted in a bad state ID error if the `nfs4_select_rw_stateid()` function returns an `-EIO` error.

BZ#1120651

A previous change to the Open vSwitch kernel module introduced a use-after-free problem that resulted in a kernel panic on systems that use this module. This update ensures that the affected object is freed on the correct place in the code, thus avoiding the problem.

BZ#1118782

Previously, the Huge Translation Lookaside Buffer (HugeTLB) unconditionally allowed access to huge pages. However, huge pages may be unsupported in some environments, such as a KVM guest on the PowerPC architecture when not backed by huge pages, and an attempt to use a base page as a huge page in memory would result in a kernel oops. This update ensures that HugeTLB denies access to huge pages if the huge pages are not supported on the system.

BZ#1096397

NFSv4 incorrectly handled a situation when an NFS client received an NFS4ERR_ADMIN_REVOKED error after sending a CLOSE operation. As a consequence, the client kept sending the same CLOSE operation indefinitely although it was receiving NFS4ERR_ADMIN_REVOKED errors. A patch has been applied to the NFSv4 code to ensure that the NFS client sends the particular CLOSE operation only once in this situation.

BZ#1099607

NFS previously called the `drop_nlink()` function after removing a file to directly decrease a link count on the related inode. Consequently, NFS did not revalidate an inode cache, and could thus use a stale file handle, resulting in an ESTALE error. A patch has been applied to ensure that NFS validates the inode cache correctly after removing a file.

BZ#1117582

A previous change to the SCSI code fixed a race condition that could occur when removing a SCSI device. However, that change caused performance degradation because it used a certain function from the block layer code that was returning different values compared with later versions of the kernel. This update alters the SCSI code to properly utilize the values returned by the block layer code.

BZ#1102794

Previously, when using a bridge interface configured on top of a bonding interface, the bonding driver was not aware of IP addresses assigned to the bridge. Consequently, with ARP monitoring enabled, the ARP monitor could not target the IP address of the bridge when probing the same subnet. The bridge was thus always reported as being down and could not be reached. With this update, the bonding driver has been made aware of IP addresses assigned to a bridge configured on top of a bonding interface, and the ARP monitor can now probe the bridge as expected. Note that the problem still occurs if the `arp_validate` option is used. Therefore, do not use this option in this case until this issue is fully resolved.

BZ#1113824

The automatic route cache rebuilding feature could incorrectly compute the length of a route hash chain if the cache contained multiple entries with the same key but a different TOS, mark, or OIF bit. Consequently, the feature could reach the rebuild limit and disable the routing cache on the system. This problem is fixed by using a helper function that avoids counting such duplicate routes.

BZ#1121541

Due to a race condition that allowed a RAID array to be written to while it was being stopped, the md driver could enter a deadlock situation. The deadlock prevented buffers from being written out to the disk, and all I/O operations to the device became unresponsive. With this update, the md driver has been modified so this deadlock is now avoided.

BZ#1112226

When booting a guest in the Hyper-V environment and enough of Programmable Interval Timer (PIT) interrupts were lost or not injected into the guest on time, the kernel panicked and the guest failed to boot. This problem has been fixed by bypassing the relevant PIT check when the guest is running under the Hyper-V environment.

All users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

8.82.4. RHSA-2014:0981 — Important: kernel security, bug fix, and enhancement update

Updated kernel packages that fix multiple security issues, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes**CVE-2014-2851, Important**

A use-after-free flaw was found in the way the `ping_init_sock()` function of the Linux kernel handled the `group_info` reference counter. A local, unprivileged user could use this flaw to crash the system or, potentially, escalate their privileges on the system.

CVE-2014-6647, Moderate

A NULL pointer dereference flaw was found in the way the `futex_wait_requeue_pi()` function of the Linux kernel's futex subsystem handled the requeuing of certain Priority Inheritance (PI) futexes. A local, unprivileged user could use this flaw to crash the system.

CVE-2014-7339, Moderate

A NULL pointer dereference flaw was found in the `rds_ib_laddr_check()` function in the Linux kernel's implementation of Reliable Datagram Sockets (RDS). A local, unprivileged user could use this flaw to crash the system.

CVE-2014-2672, Moderate

It was found that a remote attacker could use a race condition flaw in the `ath_tx_aggr_sleep()` function to crash the system by creating large network traffic on the system's Atheros 9k wireless network adapter.

CVE-2014-2678, Moderate

A NULL pointer dereference flaw was found in the `rds_iw_laddr_check()` function in the Linux kernel's implementation of Reliable Datagram Sockets (RDS). A local, unprivileged user could use this flaw to crash the system.

CVE-2014-2706, Moderate

A race condition flaw was found in the way the Linux kernel's mac80211 subsystem implementation handled synchronization between TX and STA wake-up code paths. A remote attacker could use this flaw to crash the system.

CVE-2014-3144, CVE-2014-3145, Moderate

An out-of-bounds memory access flaw was found in the Netlink Attribute extension of the Berkeley Packet Filter (BPF) interpreter functionality in the Linux kernel's networking implementation. A local, unprivileged user could use this flaw to crash the system or leak kernel memory to user space via a specially crafted socket filter.

Bug Fixes

BZ#1107503

Due to a bug in the mount option parser, prefix paths on a CIFS DFS share could be prepended with a double backslash ("\\"), resulting in an incorrect "No such file" error in certain environments. The mount option parser has been fixed and prefix paths now starts with a single backslash as expected.

BZ#1110170, BZ#1110169, BZ#1110168, BZ#1109885, BZ#1109883

Several concurrency problems, that could result in data corruption, were found in the implementation of CTR and CBC modes of operation for AES, DES, and DES3 algorithms on IBM S/390 systems. Specifically, a working page was not protected against concurrency invocation in CTR mode. The fallback solution for not getting a working page in CTR mode did not handle iv values correctly. The CBC mode used did not properly save and restore the key and iv values in some concurrency situations. All these problems have been addressed in the code and the concurrent use of the aforementioned algorithms no longer cause data corruption.

BZ#1090749

In cluster environment, the multicast traffic from the guest to a host could be sometimes unreliable. An attempt to resolve this problem was made with the RHSA-2013-1645 advisory, however, that attempt introduced a regression. This update reverts patches for this problem provided by RHSA-2013-1645 and introduces a new fix of the problem. The problem has been resolved by flooding the network with multicast packets if the multicast querier is disabled and no other querier has been detected.

BZ#1106472

The bridge MDB RTNL handlers were incorrectly removed after deleting a bridge from the system with more than one bridge configured. This led to various problems, such as that the multicast IGMP snooping data from the remaining bridges were not displayed. This update ensures that the bridge handlers are removed only after the bridge module is unloaded, and the multicast IGMP snooping data now displays correctly in the described situation.

BZ#1100574

Due to a bug in the nouveau kernel module, the wrong display output could be modified in certain multi-display configurations. Consequently, on Lenovo Thinkpad T420 and W530 laptops with an external display connected, this could result in the LVDS panel "bleeding" to white during startup, and the display controller might become non-functional until after a reboot. Changes to the display configuration could also trigger the bug under various circumstances. With this update, the nouveau kernel module has been corrected and the said configurations now work as expected.

BZ#1103821

When guest supports Supervisor Mode Execution Protection (SMEP), KVM sets the appropriate permissions bits on the guest page table entries (sptes) to emulate SMEP enforced access.

Previously, KVM was incorrectly verifying whether the "smep" bit was set in the host cr4 register instead of the guest cr4 register. Consequently, if the host supported SMEP, it was enforced even though it was not requested, which could render the guest system unbootable. This update corrects the said "smep" bit check and the guest system boot as expected in this scenario.

BZ#1096059

Previously, if a hrtimer interrupt was delayed, all future pending hrtimer events that were queued on the same processor were also delayed until the initial hrtimer event was handled. This could cause all hrtimer processing to stop for a significant period of time. To prevent this problem, the kernel has been modified to handle all expired hrtimer events when handling the initially delayed hrtimer event.

BZ#1099725

Previously, hardware could execute commands send by drivers in FIFO order instead of tagged order. Commands thus could be executed out of sequence, which could result in large latencies and degradation of throughput. With this update, the ATA subsystem tags each command sent to the hardware, ensuring that the hardware executes commands in tagged order. Performance on controllers supporting tagged commands can now increase by 30-50%.

BZ#1107931

Due to a bug in the GRE tunneling code, it was impossible to create a GRE tunnel with a custom name. This update corrects behavior of the `ip_tunnel_find()` function, allowing users to create GRE tunnels with custom names.

BZ#1110658

The `qla2xxx` driver has been upgraded to version 8.05.00.03.06.5-k2, which provides a number of bug fixes over the previous version in order to correct various timeout problems with the mailbox command.

BZ#1093984

The kernel previously did not reset the kernel ring buffer if the trace clock was changed during tracing. However, the new clock source could be inconsistent with the previous clock source, and the result trace record thus could contain incomparable time stamps. To ensure that the trace record contains only comparable time stamps, the ring buffer is now reset whenever the trace clock changes.

BZ#1103972

Previously, KVM did not accept PCI domain (segment) number for host PCI devices, making it impossible to assign a PCI device that was a part of a non-zero PCI segment to a virtual machine. To resolve this problem, KVM has been extended to accept PCI domain number in addition to slot, device, and function numbers.

Enhancement**BZ#1094403**

Users can now set ToS, TTL, and priority values in IPv4 on per-packet basis.

All users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

8.82.5. [RHSA-2014:0771](#) — Important: kernel security and bug fix update

Updated kernel packages that fix multiple security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2014-3153, Important


A flaw was found in the way the Linux kernel's futex subsystem handled the requeuing of certain Priority Inheritance (PI) futexes. A local, unprivileged user could use this flaw to escalate their privileges on the system.

CVE-2014-1737, Important

A flaw was found in the way the Linux kernel's floppy driver handled user space provided data in certain error code paths while processing FDRAWCMD IOCTL commands. A local user with write access to `/dev/fdX` could use this flaw to free (using the `kfree()` function) arbitrary kernel memory.

CVE-2014-1738, Low

It was found that the Linux kernel's floppy driver leaked internal kernel memory addresses to user space during the processing of the FDRAWCMD IOCTL command. A local user with write access to `/dev/fdX` could use this flaw to obtain information about the kernel heap arrangement.



NOTE

A local user with write access to `/dev/fdX` could use these two flaws (CVE-2014-1737 in combination with CVE-2014-1738) to escalate their privileges on the system.

CVE-2014-0203, Moderate

It was discovered that the `proc_ns_follow_link()` function did not properly return the `LAST_BIND` value in the last pathname component as is expected for `procfs` symbolic links, which could lead to excessive freeing of memory and consequent slab corruption. A local, unprivileged user could use this flaw to crash the system.

CVE-2014-2039, Moderate

A flaw was found in the way the Linux kernel handled exceptions when user-space applications attempted to use the linkage stack. On IBM S/390 systems, a local, unprivileged user could use this flaw to crash the system.

CVE-2013-6378, Low

An invalid pointer dereference flaw was found in the Marvell 8xxx Libertas WLAN (`libertas`) driver in the Linux kernel. A local user able to write to a file that is provided by the `libertas` driver and located on the debug file system (`debugfs`) could use this flaw to crash the system. Note: The `debugfs` file system must be mounted locally to exploit this issue. It is not mounted by default.

CVE-2014-1874, Low

A denial of service flaw was discovered in the way the Linux kernel's SELinux implementation handled files with an empty SELinux security context. A local user who has the CAP_MAC_ADMIN capability could use this flaw to crash the system.

Red Hat would like to thank Kees Cook of Google for reporting CVE-2014-3153, Matthew Daley for reporting CVE-2014-1737 and CVE-2014-1738, and Vladimir Davydov of Parallels for reporting CVE-2014-0203. Google acknowledges Pinkie Pie as the original reporter of CVE-2014-3153.

Bug Fixes

BZ#1086839

Due to a ndlp list corruption bug in the lpfc driver, systems with Emulex LPe16002B-M6 PCIe 2-port 16Gb Fibre Channel Adapters could trigger a kernel panic during I/O operations. A series of patches has been backported to address this problem so the kernel no longer panics during I/O operations on the aforementioned systems.

BZ#1096214

A previous change enabled receive acceleration for VLAN interfaces configured on a bridge interface. However, this change allowed VLAN-tagged packets to bypass the bridge and be delivered directly to the VLAN interfaces. This update ensures that the traffic is correctly processed by a bridge before it is passed to any VLAN interfaces configured on that bridge.

BZ#1090750

A previous change that introduced global clock updates caused guest machines to boot slowly when the host Time Stamp Counter (TSC) was marked as unstable. The slow down increased with the number of vCPUs allocated. To resolve this problem, a patch has been applied to limit the rate of the global clock updates.

BZ#1094287

Due to a bug in the ixgbev driver, the stripped VLAN information from incoming packets on the ixgbev interface could be lost, and such packets thus did not reach a related VLAN interface. This problem has been fixed by adding the packet's VLAN information to the Socket Buffer (skb) before passing it to the network stack. As a result, the ixgbev driver now passes the VLAN-tagged packets to the appropriate VLAN interface.

BZ#1089915

A race condition between completion and timeout handling in the block device code could sometimes trigger a BUG_ON() assertion, resulting in a kernel panic. This update resolves this problem by relocating a relevant function call and the BUG_ON() assertion in the code.

BZ#1088779

Systems that use NFS file systems could become unresponsive or trigger a kernel oops due to a use-after-free bug in the duplicate reply cache (DRC) code in the nfsd daemon. This problem has been resolved by modifying nfsd to unhash DRC entries before attempting to use them and to prefer to allocate a new DRC entry from the slab instead of reusing an expired entry from the list.

BZ#1092002

When an attempt to create a file on the GFS2 file system failed due to a file system quota violation, the relevant VFS inode was not completely uninitialized. This could result in a list corruption error. This update resolves this problem by correctly uninitialized the VFS inode in this situation.

BZ#1069630

Previously, automount could become unresponsive when trying to reconnect to mounts with the direct or offset mount types at system startup. This happened because the device ioctl code did not handle the situation when the relevant caller did not yet own the mount. Also, the umount() command sometimes failed to unmount an NFS file system with the stale root. Both problems have been addressed in the virtual file system code, and automount is now able to mount direct or offset mounts using a new lookup function, kern_path_mountpoint(). The umount() command now handles mount points without their revalidation, which allows the command to unmount NFS file systems with the stale root.

BZ#1091424

The kernel did not handle environmental and power warning (EPOW) interrupts correctly. This prevented successful usage of the "virsh shutdown" command to shut down guests on IBM POWER8 systems. This update ensures that the kernel handles EPOW events correctly and also prints informative descriptions for the respective EPOW events. The detailed information about each encountered EPOW can be found in the Real-Time Abstraction Service (RTAS) error log.

BZ#1081915

Due to a race condition in the cgroup code, the kernel task scheduler could trigger a kernel panic when it was moving an exiting task between cgroups. A patch has been applied to avoid this kernel panic by replacing several improperly used function calls in the cgroup code.

BZ#1081909

An incorrectly placed function call in the cgroup code prevented the notify_on_release functionality from working properly. This functionality is used to remove empty cgroup directories, however due to this bug, some empty cgroup directories were remaining on the system. This update ensures that the notify_on_release functionality is always correctly triggered by correctly ordering operations in the cgroup_task_migrate() function.

BZ#1081914

Due to a race condition in the cgroup code, the kernel task scheduler could trigger a use-after-free bug when it was moving an exiting task between cgroups, which resulted in a kernel panic. This update avoids the kernel panic by introducing a new function, cpu_cgroup_exit(). This function ensures that the kernel does not release a cgroup that is not empty yet.

BZ#1079869

Due to a bug in the hrtimers subsystem, the clock_was_set() function called an inter-processor interrupt (IPI) from soft IRQ context and waited for its completion, which could result in a deadlock situation. A patch has been applied to fix this problem by moving the clock_was_set() function call to the working context. Also during the resume process, the hrtimers_resume() function reprogrammed kernel timers only for the current CPU because it assumed that all other CPUs are offline. However, this assumption was incorrect in certain scenarios, such as when resuming a Xen guest with some non-boot CPUs being only stopped with IRQs disabled. As a consequence, kernel timers were not corrected on other than the boot CPU even though those CPUs were online. To resolve this problem, hrtimers_resume() has been modified to trigger an early soft IRQ to correctly reprogram kernel timers on all CPUs that are online.

BZ#1080104

Due to a previous change that altered the format of the txselect parameter, the InfiniBand qib driver was unable to support HP branded QLogic QDR InfiniBand cards in HP Blade servers. To resolve this problem, the driver's parsing routine, setup_txselect(), has been modified to handle multi-value strings.

BZ#1075653

A previous change to the virtual file system (VFS) code included the reduction of the `PATH_MAX` variable by 32 bytes. However, this change was not propagated to the `do_getname()` function, which had a negative impact on interactions between the `getname()` and `do_getname()` functions. This update modifies `do_getname()` accordingly and this function now works as expected.

BZ#1082622

Previously, in certain environments, such as an HP BladeSystem Enclosure with several Blade servers, the `kdump` kernel could experience a kernel panic or become unresponsive during boot due to lack of available interrupt vectors. As a consequence, `kdump` failed to capture a core dump. To increase a number of available interrupt vectors, the `kdump` kernel can boot up with more CPUs. However, the `kdump` kernel always tries to boot up with the bootstrap processor (BSP), which can cause the kernel to fail to bring up more than one CPU under certain circumstances. This update introduces a new kernel parameter, `disable_cpu_acipid`, which allows the `kdump` kernel to disable BSP during boot and then to successfully boot up with multiple processors. This resolves the problem of lack of available interrupt vectors for systems with a high number of devices and ensures that `kdump` can now successfully capture a core dump on these systems.

BZ#1091826

A previous patch to the kernel scheduler fixed a kernel panic caused by a divide-by-zero bug in the `init_numa_sched_groups_power()` function. However, that patch introduced a regression on systems with standard Non-Uniform Memory Access (NUMA) topology so that `cpu_power` in all but one NUMA domains was set to twice the expected value. This resulted in incorrect task scheduling and some processors being left idle even though there were enough queued tasks to handle, which had a negative impact on system performance. This update ensures that `cpu_power` on systems with standard NUMA topology is set to expected values by adding an estimate to `cpu_power` for every uncounted CPU. Task scheduling now works as expected on these systems without performance issues related to the said bug.

BZ#1092870

The `RTM_NEWLINK` messages can contain information about every virtual function (VF) for the given network interface (NIC) and can become very large if this information is not filtered. Previously, the kernel netlink interface allowed the `getifaddr()` function to process `RTM_NEWLINK` messages with unfiltered content. Under certain circumstances, the kernel netlink interface would omit data for the given group of NICs, causing `getifaddr()` to loop indefinitely being unable to return information about the affected NICs. This update resolves this problem by supplying only the `RTM_NEWLINK` messages with filtered content.

BZ#1063508

The `ext4_releasepage()` function previously emitted an unnecessary warning message when it was passed a page with the `PageChecked` flag set. To avoid irrelevant warnings in the kernel log, this update removes the related `WARN_ON()` from the `ext4` code.

BZ#1070296

Microsoft Windows 7 KVM guests could become unresponsive during reboot because KVM did not manage to inject a Non-Maskable Interrupt (NMI) to the guest when handling page faults. To resolve this problem, a series of patches has been applied to the KVM code, ensuring that KVM handles page faults during the reboot of the guest machine as expected.

BZ#1096711

The turbostat utility produced error messages when used on systems with the fourth generation of Intel Core Processors. To fix this problem, the kernel has been updated to provide the C-state residency information for the C8, C9, and C10 C-states.

All kernel users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

8.82.6. RHSA-2014-0475 — Important: kernel security and bug fix update

Updated kernel packages that fix three security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2014-2523, Important

A flaw was found in the way the Linux kernel's netfilter connection tracking implementation for Datagram Congestion Control Protocol (DCCP) packets used the `skb_header_pointer()` function. A remote attacker could use this flaw to send a specially crafted DCCP packet to crash the system or, potentially, escalate their privileges on the system.

CVE-2014-6383, Moderate

A flaw was found in the way the Linux kernel's Adaptec RAID controller (`aacraid`) checked permissions of `compat` IOCTLs. A local attacker could use this flaw to bypass intended security restrictions.

CVE-2014-0077, Moderate

A flaw was found in the way the `handle_rx()` function handled large network packets when mergeable buffers were disabled. A privileged guest user could use this flaw to crash the host or corrupt QEMU process memory on the host, which could potentially result in arbitrary code execution on the host with the privileges of the QEMU process.

The CVE-2014-0077 issue was discovered by Michael S. Tsirkin of Red Hat.

Bug Fixes

BZ#1078007

Due to recent changes in the Linux memory management, the kernel did not properly handle per-CPU LRU page vectors when hot unplugging CPUs. As a consequence, the page vector of the relevant offline CPU kept memory pages for memory accounting. This prevented the `libvirtd` daemon from removing the relevant memory cgroup directory upon system shutdown, rendering `libvirtd` unresponsive. To resolve this problem, the Linux memory management now properly flushes memory pages of offline CPUs from the relevant page vectors.

BZ#1063201

Recent changes in the `d_splice_alias()` function introduced a bug that allowed `d_splice_alias()` to

return a dentry from a different directory than was the directory being looked up. As a consequence in cluster environment, a kernel panic could be triggered when a directory was being removed while a concurrent cross-directory operation was performed on this directory on another cluster node. This update avoids the kernel panic in this situation by correcting the search logic in the `d_splice_alias()` function so that the function can no longer return a dentry from an incorrect directory.

BZ#1086095

A system could enter a deadlock situation when the Real-Time (RT) scheduler was moving RT tasks between CPUs and the `wakeup_kswapd()` function was called on multiple CPUs, resulting in a kernel panic. This problem has been fixed by removing a problematic memory allocation and therefore calling the `wakeup_kswapd()` function from a deadlock-safe context.

BZ#1086007

Previously some device mapper kernel modules, such as `dm-thin`, `dm-space-map-metadata`, and `dm-bufio`, contained various bugs that had adverse effects on their proper functioning. This update backports several upstream patches that resolve these problems, including a fix for the metadata resizing feature of device mapper thin provisioning (`thinp`) and fixes for read-only mode for `dm-thin` and `dm-bufio`. As a result, the aforementioned kernel modules now contain the latest upstream changes and work as expected.

BZ#1066535

A previous change in the TCP code that extended the "proto" struct with a new function, `release_cb()`, broke integrity of the kernel Application Binary Interface (kABI). If the core stack called a newly introduced pointer to this function for a module that was compiled against older kernel headers, the call resulted in out-of-bounds access and a subsequent kernel panic. To avoid this problem, the core stack has been modified to recognize a newly introduced slab flag, `RHEL_EXTENDED_PROTO`. This allows the core stack to safely access the `release_cb` pointer only for modules that support it.

BZ#1083350

The Completely Fair Scheduler (CFS) did not verify whether the CFS period timer is running while throttling tasks on the CFS run queue. Therefore under certain circumstances, the CFS run queue became stuck because the CFS period timer was inactive and could not be restarted. To fix this problem, the CFS now restarts the CFS period timer inside the throttling function if it is inactive.

BZ#1073562

A previous change removed the `ZONE_RECLAIM_LOCKED` flag from Linux memory management code in order to fix a NUMA node allocation problem in the memory zone reclaim logic. However, the flag removal allowed concurrent page reclaiming within one memory zone, which, under heavy system load, resulted in unwanted spin lock contention and subsequent performance problems (systems became slow or unresponsive). This update resolves this problem by preventing reclaim threads from scanning a memory zone if the zone does not satisfy scanning requirements. Systems under heavy load no longer suffer from CPU overloading but sustain their expected performance.

BZ#1073564

The restart logic for the memory reclaiming with compaction was previously applied on the level of LRU page vectors. This could, however, cause significant latency in memory allocation because memory compaction does not require only memory pages of a certain cgroup but a whole memory zone. This performance issue has been fixed by moving the restart logic to the zone level and restarting the memory reclaim for all memory cgroups in a zone when the compaction requires more free pages from the zone.

BZ#1074855

Previously, the `for_each_iscsi_host()` macro was incorrectly defined so it accessed an out-of-range element for a 2-element array. This macro was also wrongly optimized by GCC 4.8 so that it was executed too many times on platforms with two SCU controllers. As a consequence, the system triggered a kernel panic when entering the S3 state, or a kernel oops when removing the `iscsi` module. This update corrects the aforementioned macro and the described problems no longer occur.

BZ#1083175

A bug in the `vmxnet3` driver allowed potential race conditions to be triggered when the driver was used with the `netconsole` module. The race conditions allowed the driver's internal NAPI poll routine to run concurrently with the `netpoll` controller routine, which resulted in data corruption and a subsequent kernel panic. To fix this problem, the `vmxnet3` driver has been modified to call the appropriate interrupt handler to schedule NAPI poll requests properly.

BZ#1081908

The kernel task scheduler could trigger a race condition while migrating tasks over CPU cgroups. The race could result in accessing a task that pointed to an incorrect parent task group, causing the system to behave unpredictably, for example to appear being unresponsive. This problem has been resolved by ensuring that the correct task group information is properly stored during the task's migration.

BZ#1076056

A previously backported patch to the XFS code added an unconditional call to the `xlog_cil_empty()` function. If the XFS file system was mounted with the unsupported `nodelaylog` option, that call resulted in access to an uninitialized spin lock and a consequent kernel panic. To avoid this problem, the `nodelaylog` option has been disabled; the option is still accepted but has no longer any effect. (The `nodelaylog` mount option was originally intended only as a testing option upstream, and has since been removed.)

BZ#1076242

The SCTP `sctp_connectx()` ABI did not work properly for 64-bit kernels compiled with 32-bit emulation. As a consequence, applications utilizing the `sctp_connectx()` function did not run in this case. To fix this problem, a new ABI has been implemented; the COMPAT ABI enables to copy and transform user data from a COMPAT-specific structure to a SCTP-specific structure. Applications that require `sctp_connectx()` now work without any problems on a system with a 64-bit kernel compiled with 32-bit emulation.

BZ#1085660

A bug in the `qla2xxx` driver caused the kernel to crash. This update resolves this problem by fixing an incorrect condition in the "for" statement in the `qla2x00_alloc_ioCBS()` function.

BZ#1079870

The code responsible for creating and binding of packet sockets was not optimized and therefore applications that utilized the `socket()` and `bind()` system calls did not perform as expected. A patch has been applied to the packet socket code so that latency for socket creating and binding is now significantly lower in certain cases.

BZ#1077874

Previously, the `vmw_pvscsi` driver could attempt to complete a command to the SCSI mid-layer after reporting a successful abort of the command. This led to a double completion bug and a subsequent kernel panic. This update ensures that the `pvscsi_abort()` function returns SUCCESS only after the abort is completed, preventing the driver from invalid attempts to complete the command.

BZ#1085658

Due to a bug in the `mlx4_en` module, a data structure related to time stamping could be accessed before being initialized. As a consequence, loading `mlx4_en` could result in a kernel crash. This problem has been fixed by moving the initiation of the time stamp mechanism to the correct place in the code.

BZ#1078011

Due to a previous change that was refactoring the Generic Routing Encapsulation (GRE) tunneling code, the `ip_gre` module did not work properly. As a consequence, GRE interfaces dropped every packet that had the Explicit Congestion Notification (ECN) bit set and did not have the ECN-Capable Transport (ECT) bit set. This update reintroduces the `ipgre_ecn_decapsulate()` function that is now used instead of the `IP_ECN_decapsulate()` function that was not properly implemented. The `ip_gre` module now works correctly and GRE devices process all packets as expected.

BZ#1078641

A bug in the `megaraid_sas` driver could cause the driver to read the hardware status values incorrectly. As a consequence, the RAID card was disabled during the system boot and the system could fail to boot. With this update, the `megaraid_sas` driver has been corrected to enable the RAID card on system boot as expected.

BZ#1081907

A bug in the Completely Fair Scheduler (CFS) could, under certain circumstances, trigger a race condition while moving a forking task between cgroups. This race could lead to a free-after-use error and a subsequent kernel panic when a child task was accessed while it was pointing to a stale cgroup of its parent task. A patch has been applied to the CFS to ensure that a child task always points to the valid parent's task group.

BZ#1078874

The Red Hat GFS2 file system previously limited a number of ACL entries per inode to 25. However, this number was insufficient in some cases, causing the `setfacl` command to fail. This update increases this limit to maximum of 300 ACL entries for the 4 KB block size. If the block size is smaller, this value is adjusted accordingly.

BZ#1085358

Previous patches to the CIFS code introduced a regression that prevented users from mounting a CIFS share using the NetBIOS over TCP service on the port 139. This problem has been fixed by masking off the top byte in the `get_rfc1002_length()` function.

BZ#1079872

Previously, user space packet capturing libraries, such as `libcap`, had a limited possibility to determine which Berkeley Packet Filter (BPF) extensions are supported by the current kernel. This limitation had a negative effect on VLAN packet filtering that is performed by the `tcpdump` utility and `tcpdump` sometimes was not able to capture filtered packets correctly. Therefore, this update introduces a new option, `SO_BPF_EXTENSIONS`, which can be specified as an argument of the `getsockopt()` function. This option enables packet capturing tools to obtain information about which BPF extensions are supported by the current kernel. As a result, the `tcpdump` utility can now capture packets properly.

BZ#1080600

The `iscsi` driver previously triggered an erroneous `BUG_ON()` assertion in case of a hard reset timeout in the `sci_apc_agent_link_up()` function. If a SATA device was unable to restore the link in time after

the reset, the iscsi port had to return to the "awaiting link-up" state. However in such a case, the port may not have been in the "resetting" state, causing a kernel panic. This problem has been fixed by removing that incorrect `BUG_ON()` assertion.

BZ#1078798

Previously, when removing an IPv6 address from an interface, unreachable routes related to that address were not removed from the IPv6 routing table. This happened because the IPv6 code used inappropriate function when searching for the routes. To avoid this problem, the IPv6 code has been modified to use the `ip6_route_lookup()` function instead of `rt6_lookup()` in this situation. All related routes are now properly deleted from the routing tables when an IPv6 address is removed.

BZ#1075651

If the BIOS returned a negative value for the critical trip point for the given thermal zone during a system boot, the whole thermal zone was invalidated and an ACPI error was printed. However, the thermal zone may still have been needed for cooling. With this update, the ACPI thermal management has been modified to only disable the relevant critical trip point in this situation.

BZ#1075554

When allocating kernel memory, the SCSI device handlers called the `sizeof()` function with a structure name as its argument. However, the modified files were using an incorrect structure name, which resulted in an insufficient amount of memory being allocated and subsequent memory corruption. This update modifies the relevant `sizeof()` function calls to rather use a pointer to the structure instead of the structure name so that the memory is now always allocated correctly.

BZ#1069848

A previous change that modified the `linkat()` system call introduced a mount point reference leak and a subsequent memory leak in case that a file system link operation returned the `ESTALE` error code. These problems have been fixed by properly freeing the old mount point reference in such a case.

BZ#1086490

The `dm-bufio` driver did not call the `blk_unplug()` function to flush plugged I/O requests. Therefore, the requests submitted by `dm-bufio` were delayed by 3 ms, which could cause performance degradation. With this update, `dm-bufio` calls `blk_unplug()` as expected, avoiding any related performance issues.

All kernel users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

8.82.7. [RHSA-2014:0328](#) — Important: kernel security and bug fix update

Updated kernel packages that fix multiple security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links in the References section.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

[CVE-2014-0055](#), Important

A flaw was found in the way the `get_rx_bufs()` function in the `vhost_net` implementation in the Linux kernel handled error conditions reported by the `vhost_get_vq_desc()` function. A privileged guest user could use this flaw to crash the host.

CVE-2014-0101, Important

A flaw was found in the way the Linux kernel processed an authenticated `COOKIE_ECHO` chunk during the initialization of an SCTP connection. A remote attacker could use this flaw to crash the system by initiating a specially crafted SCTP handshake in order to trigger a NULL pointer dereference on the system.

CVE-2014-0069, Moderate

A flaw was found in the way the Linux kernel's CIFS implementation handled uncached write operations with specially crafted `iovec` structures. An unprivileged local user with access to a CIFS share could use this flaw to crash the system, leak kernel memory, or, potentially, escalate their privileges on the system. Note: the default cache settings for CIFS mounts on Red Hat Enterprise Linux 6 prohibit a successful exploitation of this issue.

CVE-2013-1860, Low

A heap-based buffer overflow flaw was found in the Linux kernel's `cdc-wdm` driver, used for USB CDC WCM device management. An attacker with physical access to a system could use this flaw to cause a denial of service or, potentially, escalate their privileges.

Red Hat would like to thank Nokia Siemens Networks for reporting CVE-2014-0101, and Al Viro for reporting CVE-2014-0069.

Bug Fixes

BZ#1063507

A previous change in the Advanced Programmable Interrupt Controller (APIC) code caused a regression on certain Intel CPUs using a Multiprocessor (MP) table. An attempt to read from the local APIC (LAPIC) could be performed before the LAPIC was mapped, resulting in a kernel crash during a system boot. A patch has been applied to fix this problem by mapping the LAPIC as soon as possible when parsing the MP table.

BZ#1067775

When removing an inode from a name space on an XFS file system, the file system could enter a deadlock situation and become unresponsive. This happened because the removal operation incorrectly used the AGF and AGI locks in the opposite order than was required by the ordering constraint, which led to a possible deadlock between the file removal and inode allocation and freeing operations. With this update, the inode's reference count is dropped before removing the inode entry with the first transaction of the removal operation. This ensures that the AGI and AGF locks are locked in the correct order, preventing any further deadlocks in this scenario.

BZ#1064913

Previously, the GFS2 kernel module leaked memory in the `gfs2_bufdata` slab cache and allowed a use-after-free race condition to be triggered in the `gfs2_remove_from_journal()` function. As a consequence after unmounting the GFS2 file system, the GFS2 slab cache could still contain some objects, which subsequently could, under certain circumstances, result in a kernel panic. A series of patches has been applied to the GFS2 kernel module, ensuring that all objects are freed from the slab cache properly and the kernel panic is avoided.

BZ#1054072

Due to the locking mechanism that the kernel used while handling Out of Memory (OOM) situations in memory control groups (cgroups), the OOM killer did not work as intended in case that many processes triggered an OOM. As a consequence, the entire system could become or appear to be unresponsive. A series of patches has been applied to improve this locking mechanism so that the OOM killer now works as expected in memory cgroups under heavy OOM load.

BZ#1055364

Previously, certain SELinux functions did not correctly handle the TCP synchronize-acknowledgment (SYN-ACK) packets when processing IPv4 labeled traffic over an INET socket. The initial SYN-ACK packets were labeled incorrectly by SELinux, and as a result, the access control decision was made using the server socket's label instead of the new connection's label. In addition, SELinux was not properly inspecting outbound labeled IPsec traffic, which led to similar problems with incorrect access control decisions. A series of patches that addresses these problems has been applied to SELinux. The initial SYN-ACK packets are now labeled correctly and SELinux processes all SYN-ACK packets as expected.

BZ#1063199

In Red Hat Enterprise Linux 6.5, the TCP Segmentation Offload (TSO) feature is automatically disabled if the corresponding network device does not report any CSUM flag in the list of its features. Previously, VLAN devices that were configured over bonding devices did not propagate its NETIF_F_NO_CSUM flag as expected, and their feature lists thus did not contain any CSUM flags. As a consequence, the TSO feature was disabled for these VLAN devices, which led to poor bandwidth performance. With this update, the bonding driver propagates the aforementioned flag correctly so that network traffic now flows through VLAN devices over bonding without any performance problems.

BZ#1064464

Due to a bug in the Infiniband driver, the ip and ifconfig utilities reported the link status of the IP over Infiniband (IPoIB) interfaces incorrectly (as "RUNNING" in case of "ifconfig", and as "UP" in case of "ip") even if no cable was connected to the respective network card. The problem has been corrected by calling the respective netif_carrier_off() function on the right place in the code. The link status of the IPoIB interfaces is now reported correctly in the described situation.

BZ#1058418

When performing read operations on an XFS file system, failed buffer readahead can leave the buffer in the cache memory marked with an error. This could lead to incorrect detection of stale errors during completion of an I/O operation because most callers do not zero out the b_error field of the buffer on a subsequent read. To avoid this problem and ensure correct I/O error detection, the b_error field of the used buffer is now zeroed out before submitting an I/O operation on a file.

BZ#1062113

Previously, when hot adding memory to the system, the memory management subsystem always performed unconditional page-block scans for all memory sections being set online. The total duration of the hot add operation depends on both, the size of memory that the system already has and the size of memory that is being added. Therefore, the hot add operation took an excessive amount of time to complete if a large amount of memory was added or if the target node already had a considerable amount of memory. This update optimizes the code so that page-block scans are performed only when necessary, which greatly reduces the duration of the hot add operation.

BZ#1059991

Due to a bug in the SELinux socket receive hook, network traffic was not dropped upon receiving a

peer:recv access control denial on some configurations. A broken labeled networking check in the SELinux socket receive hook has been corrected, and network traffic is now properly dropped in the described case.

BZ#1060491

When transferring a large amount of data over the peer-to-peer (PPP) link, a rare race condition between the throttle() and unthrottle() functions in the tty driver could be triggered. As a consequence, the tty driver became unresponsive, remaining in the throttled state, which resulted in the traffic being stalled. Also, if the PPP link was heavily loaded, another race condition in the tty driver could have been triggered. This race allowed an unsafe update of the available buffer space, which could also result in the stalled traffic. A series of patches addressing both race conditions has been applied to the tty driver; if the first race is triggered, the driver loops and forces re-evaluation of the respective test condition, which ensures uninterrupted traffic flow in the described situation. The second race is now completely avoided due to a well-placed read lock, and the update of the available buffer space proceeds correctly.

BZ#1058420

Previously, the e752x_edac module incorrectly handled the pci_dev usage count, which could reach zero and deallocate a PCI device structure. As a consequence, a kernel panic could occur when the module was loaded multiple times on some systems. This update fixes the usage count that is triggered by loading and unloading the module repeatedly, and a kernel panic no longer occurs.

BZ#1057165

When a page table is upgraded, a new top level of the page table is added for the virtual address space, which results in a new Address Space Control Element (ASCE). However, the Translation Lookaside Buffer (TLB) of the virtual address space was not previously flushed on page table upgrade. As a consequence, the TLB contained entries associated with the old ASCE, which led to unexpected program failures and random data corruption. To correct this problem, the TLB entries associated with the old ASCE are now flushed as expected upon page table upgrade.

BZ#1064115

When a network interface is running in promiscuous (PROMISC) mode, the interface may receive and process VLAN-tagged frames even though no VLAN is attached to the interface. However, the enic driver did not handle processing of the packets with the VLAN-tagged frames in PROMISC mode correctly if the frames had no VLAN group assigned, which led to various problems. To handle the VLAN-tagged frames without a VLAN group properly, the frames have to be processed by the VLAN code, and the enic driver thus no longer verifies whether the packet's VLAN group field is empty.

BZ#1057164

A previous change in the Linux memory management on IBM System z removed the handler for the Address Space Control Element (ASCE) type of exception. As a consequence, the kernel was unable to handle ASCE exceptions, which led to a kernel panic. Such an exception was triggered, for example, if the kernel attempted to access user memory with an address that was larger than the current page table limit from a user-space program. This problem has been fixed by calling the standard page fault handler, do_dat_exception, if an ASCE exception is raised.

BZ#1063271

Due to several bugs in the network console logging, a race condition between the network console send operation and the driver's IRQ handler could occur, or the network console could access invalid memory content. As a consequence, the respective driver, such as vmxnet3, triggered a BUG_ON() assertion and the system terminated unexpectedly. A patch addressing these bugs has been applied

so that driver's IRQs are disabled before processing the send operation and the network console now accesses the RCU-protected (read-copy update) data properly. Systems using the network console logging no longer crashes due to the aforementioned conditions.

All kernel users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

8.82.8. **RHSA-2014:0159 — Important: kernel security and bug fix update**

Updated kernel packages that fix multiple security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having Important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2013-6381, Important

A buffer overflow flaw was found in the way the `qeth_snmp_command()` function in the Linux kernel's QETH network device driver implementation handled SNMP IOCTL requests with an out-of-bounds length. A local, unprivileged user could use this flaw to crash the system or, potentially, escalate their privileges on the system.

CVE-2013-2929, Low

A flaw was found in the way the `get_dumpable()` function return value was interpreted in the `ptrace` subsystem of the Linux kernel. When `'fs.suid_dumpable'` was set to 2, a local, unprivileged local user could use this flaw to bypass intended `ptrace` restrictions and obtain potentially sensitive information.

CVE-2013-7263, CVE-2013-7265, Low

It was found that certain protocol handlers in the Linux kernel's networking implementation could set the `addr_len` value without initializing the associated data structure. A local, unprivileged user could use this flaw to leak kernel stack memory to user space using the `recvmsg`, `recvfrom`, and `recvmsg` system calls.

Bug Fixes

BZ#1051393

Due to a bug in the NFS code, the state manager and the `DELEGRETURN` operation could enter a deadlock if an asynchronous session error was received while `DELEGRETURN` was being processed by the state manager. The state manager became unable to process the failing `DELEGRETURN` operation because it was waiting for an asynchronous RPC task to complete, which could not have been completed because the `DELEGRETURN` operation was cycling indefinitely with session errors. A series of patches has been applied to ensure that the asynchronous error handler waits for recovery when a session error is received and the deadlock no longer occurs.

BZ#1049590

The IPv4 and IPv6 code contained several issues related to the `conntrack` fragmentation handling that prevented fragmented packages from being properly reassembled. This update applies a series

of patches and ensures that MTU discovery is handled properly, and fragments are correctly matched and packets reassembled.

BZ#1046043

Inefficient usage of Big Kernel Locks (BKLs) in the `ptrace()` system call could lead to BKL contention on certain systems that widely utilize `ptrace()`, such as User-mode Linux (UML) systems, resulting in degraded performance on these systems. This update removes the relevant BKLs from the `ptrace()` system call, thus resolving any related performance issues.

BZ#1046041

When utilizing SCTP over the bonding device in Red Hat Enterprise Linux 6.5, SCTP assumed offload capabilities on virtual devices where it was not guaranteed that underlying physical devices are equipped with these capabilities. As a consequence, checksums of the outgoing packets became corrupted and a network connection could not be properly established. A patch has been applied to ensure that checksums of the packages to the devices without SCTP checksum capabilities are properly calculated in software fallback. SCTP connections over the bonding devices can now be established as expected in Red Hat Enterprise Linux 6.5.

BZ#1044566

The context of the user's process could not be previously saved on PowerPC platforms if the VSX Machine State Register (MSR) bit was set but the user did not provide enough space to save the VSX state. This update allows to clear the VSX MSR bit in such a situation, indicating that there is no valid VSX state in the user context.

BZ#1043779

After a statically defined gateway became unreachable and its corresponding neighbor entry entered a FAILED state, the gateway stayed in the FAILED state even after it became reachable again. As a consequence, traffic was not routed through that gateway. This update enables probing such a gateway automatically so that the traffic can be routed through this gateway again once it becomes reachable.

BZ#1040826

Due to several bugs in the IPv6 code, a soft lockup could occur when the number of cached IPv6 destination entries reached the garbage collector threshold on a high-traffic router. A series of patches has been applied to address this problem. These patches ensure that the route probing is performed asynchronously to prevent a dead lock with garbage collection. Also, the garbage collector is now run asynchronously, preventing CPUs that concurrently requested the garbage collector from waiting until all other CPUs finish the garbage collection. As a result, soft lockups no longer occur in the described situation.

BZ#1035347

A previous change to the md driver disabled the TRIM operation for RAID5 volumes in order to prevent a possible kernel oops. However, if a MD RAID volume was reshaped to a different RAID level, this could result in TRIM being disabled on the resulting volume, as the RAID4 personality is used for certain reshapes. A patch has been applied that corrects this problem by setting the stacking limits before changing a RAID level, and thus ensuring the correct discard (TRIM) granularity for the RAID array.

BZ#1051395

NFS previously allowed a race between "silly rename" operations and the `rmdir()` function to occur when removing a directory right after an unlinked file in the directory was closed. As a result, `rmdir()` could fail with an EBUSY error. This update applies a patch ensuring that NFS waits for any

asynchronous operations to complete before performing the `rmdir()` operation.

BZ#1051394

Due to a bug in the EDAC driver, the driver failed to decode and report errors on AMD family 16h processors correctly. This update incorporates a missing case statement to the code so that the EDAC driver now handles errors as expected.

BZ#1045094

A deadlock between the state manager, `kswapd` daemon, and the `sys_open()` function could occur when the state manager was recovering from an expired state and recovery OPEN operations were being processed. To fix this problem, NFS has been modified to ignore all errors from the LAYOUTRETURN operation (a pNFS operation) except for "NFS4ERR_DELAY" in this situation.

BZ#1040498

The `bnx2x` driver handled unsupported TLVs received from a Virtual Function (VF) using the VF-PF channel incorrectly; when a driver of the VF sent a known but unsupported TLV command to the Physical Function, the driver of the PF did not reply. As a consequence, the VF-PF channel was left in an unstable state and the VF eventually timed out. A patch has been applied to correct the VF-PF locking scheme so that unsupported TLVs are properly handled and responded to by the PF side. Also, unsupported TLVs could previously render a mutex used to lock the VF-PF operations. The mutex then stopped protecting critical sections of the code, which could result in error messages being generated when the PF received additional TLVs from the VF. A patch has been applied that corrects the VF-PF channel locking scheme, and unsupported TLVs thus can no longer break the VF-PF lock.

BZ#1040497

A bug in the statistics flow in the `bnx2x` driver caused the card's DMA Engine (DMAE) to be accessed without taking a necessary lock. As a consequence, previously queued DMAE commands could be overwritten and the Virtual Functions then could timeout on requests to their respective Physical Functions. The likelihood of triggering the bug was higher with more SR-IOV Virtual Functions configured. Overwriting of the DMAE commands could also result in other problems even without using SR-IOV. This update ensures that all flows utilizing DMAE will use the same API and the proper locking scheme is kept by all these flows.

BZ#1035339

When starting or waking up a system that utilized an AHCI controller with empty ports, and the EM transmit bit was busy, the AHCI driver incorrectly released the related error handler before initiation of the sleep operation. As a consequence, the error handler could be acquired by a different port of the AHCI controller and the Serial General Purpose Input/Output (SGPIO) signal could eventually blink the rebuild pattern on an empty port. This update implements cross-port error handler exclusion to the generic ATA driver and the AHCI driver has been modified to use the `msleep()` function in this particular case. The error handler is no longer released upon the sleep operation and the SGPIO signal can no longer indicate the disk's rebuild on the empty controller's slot.

BZ#1032389

Previous changes to the `igb` driver caused the `ethtool` utility to determine and display some capabilities of the Ethernet devices incorrectly. This update fixes the `igb` driver so that the actual link capabilities are now determined properly, and `ethtool` displays values as accurate as possible in dependency on the data available to the driver.

All users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

8.82.9. **RHSA-2013:1801 — Important: kernel security and bug fix update**

Updated kernel packages that fix multiple security issues, several bugs, and add two enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2013-4470, Important

A flaw was found in the way the Linux kernel's TCP/IP protocol suite implementation handled sending of certain UDP packets over sockets that used the UDP_CORK option when the UDP Fragmentation Offload (UFO) feature was enabled on the output device. A local, unprivileged user could use this flaw to cause a denial of service or, potentially, escalate their privileges on the system.

CVE-2013-6367, Important

A divide-by-zero flaw was found in the `apic_get_tmctt()` function in KVM's Local Advanced Programmable Interrupt Controller (LAPIC) implementation. A privileged guest user could use this flaw to crash the host.

CVE-2013-6368, Important

A memory corruption flaw was discovered in the way KVM handled virtual APIC accesses that crossed a page boundary. A local, unprivileged user could use this flaw to crash the system or, potentially, escalate their privileges on the system.

CVE-2013-2141, Low

An information leak flaw in the Linux kernel could allow a local, unprivileged user to leak kernel memory to user space.

Red Hat would like to thank Hannes Frederic Sowa for reporting CVE-2013-4470, and Andrew Honig of Google for reporting CVE-2013-6367 and CVE-2013-6368.

Bug Fixes

BZ#1027343

Due to a regression bug in the `mlx4` driver, Mellanox `mlx4` adapters could become unresponsive on heavy load along with IOMMU allocation errors being logged to the systems logs. A patch has been applied to the `mlx4` driver so that the driver now calculates the last memory page fragment when allocating memory in the Rx path.

BZ#1028278

A bug in the RSXX DMA handling code allowed DISCARD operations to call the `pci_unmap_page()` function, which triggered a race condition on the PowerPC architecture when DISCARD, READ, and WRITE operations were issued simultaneously. However, DISCARD operations are always assigned

a DMA address of 0 because they are never mapped. Therefore, this race could result in freeing memory that was mapped for another operation and a subsequent EEH event. A patch has been applied, preventing the DISCARD operations from calling `pci_unmap_page()`, and thus avoiding the aforementioned race condition.

BZ#1029330

Due to a missing part of the bcma driver, the brcmsmac kernel module did not have a list of internal aliases that was needed by the kernel to properly handle the related udev events. Consequently, when the bcma driver scanned for the devices at boot time, these udev events were ignored and the kernel did not load the brcmsmac module automatically. A patch that provides missing aliases has been applied so that the udev requests of the brcmsmac module are now handled as expected and the kernel loads the brcmsmac module automatically on boot.

BZ#1029997

A bug in the mlx4 driver could trigger a race between the "blue flame" feature's traffic flow and the stamping mechanism in the Tx ring flow when processing Work Queue Elements (WQEs) in the Tx ring. Consequently, the related queue pair (QP) of the mlx4 Ethernet card entered an error state and the traffic on the related Tx ring was blocked. A patch has been applied to the mlx4 driver so that the driver does not stamp the last completed WQE in the Tx ring, and thus avoids the aforementioned race.

BZ#1030171

A previous change in the NFSv4 code resulted in breaking the sync NFSv4 mount option. A patch has been applied that restores functionality of the sync mount option.

BZ#1030713

Due to a bug in the Emulex lpfc driver, the driver could not allocate a SCSI buffer properly, which resulted in severe performance degradation of lpfc adapters on 64-bit PowerPC systems. A patch addressing this problem has been applied so that lpfc allocates the SCSI buffer correctly and lpfc adapters now work as expected on 64-bit PowerPC systems.

BZ#1032162

When performing I/O operations on a heavily-fragmented GFS2 file system, significant performance degradation could occur. This was caused by the allocation strategy that GFS2 used to search for an ideal contiguous chunk of free blocks in all the available resource groups (rgrp). A series of patches has been applied that improves performance of GFS2 file systems in case of heavy fragmentation. GFS2 now allocates the biggest extent found in the rgrp if it fulfills the minimum requirements. GFS2 has also reduced the amount of bitmap searching in case of multi-block reservations by keeping track of the smallest extent for which the multi-block reservation would fail in the given rgrp. This improves GFS2 performance by avoiding unnecessary rgrp free block searches that would fail. Additionally, this patch series fixes a bug in the GFS2 block allocation code where a multi-block reservation was not properly removed from the rgrp's reservation tree when it was disqualified, which eventually triggered a `BUG_ON()` macro due to an incorrect count of reserved blocks.

BZ#1032167

An earlier patch to the kernel added the dynamic queue depth throttling functionality to the QLogic's qla2xxx driver that allowed the driver to adjust queue depth for attached SCSI devices. However, the kernel might have crashed when having this functionality enabled in certain environments, such as on systems with EMC PowerPath Multipathing installed that were under heavy I/O load. To resolve this problem, the dynamic queue depth throttling functionality has been removed from the qla2xxx driver.

BZ#1032168

Previously, devices using the ixgbev driver that were assigned to a virtual machine could not adjust their Jumbo MTU value automatically if the Physical Function (PF) interface was down; when the PF device was brought up, the MTU value on the related Virtual Function (VF) device was set incorrectly. This was caused by the way the communication channel between PF and VF interfaces was set up and the first negotiation attempt between PF and VF was made. To fix this problem, structural changes to the ixgbev driver have been made so that the kernel can now negotiate the correct API between PF and VF successfully and the MTU value is now set correctly on the VF interface in this situation.

BZ#1032170

A bug in the ixgbe driver caused that IPv6 hardware filtering tables were not correctly rewritten upon interface reset when using a bridge device over the PF interface in an SR-IOV environment. As a result, the IPv6 traffic between VFs was interrupted. An upstream patch has been backported to modify the ixgbe driver so that the update of the Multimedia Terminal Adapter (MTA) table is now unconditional, avoiding possible inconsistencies in the MTA table upon PF's reset. The IPv6 traffic between VFs proceeds as expected in this scenario.

BZ#1032247

When using Haswell HDMI audio controllers with an unaligned DMA buffer size, these audio controllers could become locked up until the next reboot for certain audio stream configurations. A patch has been applied to the Intel's High Definition Audio (HDA) driver that enforces the DMA buffer alignment setting for the Haswell HDMI audio controllers. These audio controllers now work as expected.

BZ#1032249

As a result of a recent fix preventing a deadlock upon an attempt to cover an active XFS log, the behavior of the `xfs_log_need_covered()` function has changed. However, `xfs_log_need_covered()` is also called to ensure that the XFS log tail is correctly updated as a part of the XFS journal sync operation. As a consequence, when shutting down an XFS file system, the sync operation failed and some files might have been lost. A patch has been applied to ensure that the tail of the XFS log is updated by logging a dummy record to the XFS journal. The sync operation completes successfully and files are properly written to the disk in this situation.

BZ#1032250

A chunk of a patch was left out when backporting a batch of patches that fixed an infinite loop problem in the LOCK operation with zero state ID during NFSv4 state ID recovery. As a consequence, the system could become unresponsive on numerous occasions. The missing chunk of the patch has been added, resolving this hang issue.

BZ#1032260

When performing buffered WRITE operations from multiple processes to a single file, the NFS code previously always verified whether the lock owner information is identical for the file being accessed even though no file locks were involved. This led to performance degradation because forked child processes had to synchronize dirty data written to a disk by the parent process before writing to a file. Also, when coalescing requests into a single READ or WRITE RPC call, NFS refused the request if the lock owner information did not match for the given file even though no file locks were involved. This also caused performance degradation. A series of patches has been applied that relax relevant test conditions so that lock owner compatibility is no longer verified in the described cases, which resolves these performance issues.

BZ#1032395

Due to a bug in the mlx4 driver, Mellanox Ethernet cards were brought down unexpectedly while

adjusting their Tx or Rx ring. A patch has been applied so that the mlx4 driver now properly verifies the state of the Ethernet card when the coalescing of the Tx or Rx ring is being set, which resolves this problem.

BZ#1032423

When the system was under memory stress, a double-free bug in the tg3 driver could have been triggered, resulting in a NIC being brought down unexpectedly followed by a kernel panic. A patch has been applied that restructures the respective code so that the affected ring buffer is freed correctly.

BZ#1032424

The RPC client always retransmitted zero-copy of the page data if it timed out before the first RPC transmission completed. However, such a retransmission could cause data corruption if using the O_DIRECT buffer and the first RPC call completed while the respective TCP socket still held a reference to the pages. To prevent the data corruption, retransmission of the RPC call is, in this situation, performed using the sendmsg() function. The sendmsg() function retransmits an authentic reproduction of the first RPC transmission because the TCP socket holds the full copy of the page data.

BZ#1032688

When creating an XFS file system, an attempt to cover an active XFS log could, under certain circumstances, result in a deadlock between the xfssyncd and xfsbufd daemons. Consequently, several kernel threads became unresponsive and the XFS file system could not have been successfully created, leading to a kernel oops. A patch has been applied to prevent this situation by forcing the active XFS log onto a disk.

Enhancements**BZ#1020518**

The kernel now supports memory configurations with more than 1TB of RAM on AMD systems.

BZ#1032426

The kernel has been modified to stop reporting ABS_MISC events on Wacom touch devices in order to ensure that the devices are correctly recognized by the HAL daemon.

All kernel users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements. The system must be rebooted for this update to take effect.

8.82.10. RHSA-2013:1645 — Important: Red Hat Enterprise Linux 6 kernel update

Updated kernel packages that fix multiple security issues, address several hundred bugs, and add numerous enhancements are now available as part of the ongoing support and maintenance of Red Hat Enterprise Linux version 6. This is the fifth regular update.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2013-4387, Important

A flaw was found in the way the Linux kernel's IPv6 implementation handled certain UDP packets when the UDP Fragmentation Offload (UFO) feature was enabled. A remote attacker could use this flaw to crash the system or, potentially, escalate their privileges on the system.

CVE-2013-0343, Moderate

A flaw was found in the way the Linux kernel handled the creation of temporary IPv6 addresses. If the IPv6 privacy extension was enabled (`/proc/sys/net/ipv6/conf/eth0/use_tempaddr` set to '2'), an attacker on the local network could disable IPv6 temporary address generation, leading to a potential information disclosure.

CVE-2013-2888, Moderate

A flaw was found in the way the Linux kernel handled HID (Human Interface Device) reports with an out-of-bounds Report ID. An attacker with physical access to the system could use this flaw to crash the system or, potentially, escalate their privileges on the system.

CVE-2013-4345, Moderate

An off-by-one flaw was found in the way the ANSI CPRNG implementation in the Linux kernel processed non-block size aligned requests. This could lead to random numbers being generated with less bits of entropy than expected when ANSI CPRNG was used.

CVE-2013-4591, Moderate

It was found that the fix for CVE-2012-2375 released via RHSA-2012:1580 accidentally removed a check for small-sized result buffers. A local, unprivileged user with access to an NFSv4 mount with ACL support could use this flaw to crash the system or, potentially, escalate their privileges on the system .

CVE-2013-4592, Moderate

A flaw was found in the way IOMMU memory mappings were handled when moving memory slots. A malicious user on a KVM host who has the ability to assign a device to a guest could use this flaw to crash the host.

CVE-2013-2889, CVE-2013-2892, Moderate

Heap-based buffer overflow flaws were found in the way the Zeroplus and Pantherlord/GreenAsia game controllers handled HID reports. An attacker with physical access to the system could use these flaws to crash the system or, potentially, escalate their privileges on the system.

CVE-2012-6542, CVE-2013-3231, Low

Two information leak flaws were found in the logical link control (LLC) implementation in the Linux kernel. A local, unprivileged user could use these flaws to leak kernel stack memory to user space.

CVE-2013-1929, Low

A heap-based buffer overflow in the way the tg3 Ethernet driver parsed the vital product data (VPD) of devices could allow an attacker with physical access to a system to cause a denial of service or, potentially, escalate their privileges.

CVE-2012-6545, CVE-2013-1928, CVE-2013-2164, CVE-2013-2234, Low

Information leak flaws in the Linux kernel could allow a privileged, local user to leak kernel memory to user space.

CVE-2013-2851, Low

A format string flaw was found in the Linux kernel's block layer. A privileged, local user could potentially use this flaw to escalate their privileges to kernel level (ring0).

Red Hat would like to thank Stephan Mueller for reporting CVE-2013-4345, and Kees Cook for reporting CVE-2013-2851.

Bug Fixes

BZ#955712

A function in the RPC code responsible for verifying whether the cached credentials matches the current process did not perform the check correctly. The code checked only whether the groups in the current process credentials appear in the same order as in the cached credential but did not ensure that no other groups are present in the cached credentials. As a consequence, when accessing files in NFS mounts, a process with the same UID and GID as the original process but with a non-matching group list could have been granted an unauthorized access to a file, or under certain circumstances, the process could have been wrongly prevented from accessing the file. The incorrect test condition has been fixed and the problem can no longer occur.

BZ#629857

When the state of the netfilter module was out-of-sync, a TCP connection was recorded in the conntrack table although the TCP connection did not exist between two hosts. If a host re-established this connection with the same source, port, destination port, source address and destination address, the host sent a TCP SYN packet and the peer sent back acknowledgment for this SYN package. However, because netfilter was out-of-sync, netfilter dropped this acknowledgment, and deleted the connection item from the conntrack table, which consequently caused the host to retransmit the SYN packet. A patch has been applied to improve this handling; if an unexpected SYN packet appears, the TCP options are annotated. Acknowledgment for the SYN packet serves as a confirmation of the connection tracking being out-of-sync, then a new connection record is created using the information annotated previously to avoid the retransmission delay.

BZ#955807

Due to several bugs in the ext4 code, data integrity system calls did not always properly persist data on the disk. Therefore, the unsynchronized data in the ext4 file system could have been lost after the system's unexpected termination. A series of patches has been applied to the ext4 code to address this problem, including a fix that ensures proper usage of data barriers in the code responsible for file synchronization. Data loss no longer occurs in the described situation.

BZ#953630

C-states for the Intel Family 6, Model 58 and 62, processors were not properly initialized in Red Hat Enterprise Linux 6. Consequently, these processors were unable to enter deep C-states. Also, C-state accounting was not functioning properly and power management tools, such as powertop or turbostat, thus displayed incorrect C-state transitions. This update applies a patch that ensures proper C-states initialization so the aforementioned processors can now enter deep core power states as expected. Note that this update does not correct C-state accounting which has been addressed by a separate patch.

BZ#953342

The kernel previously did not handle situation where the system needed to fall back from non-flat

Advanced Programmable Interrupt Controller (APIC) mode to flat APIC mode. Consequently, a NULL pointer was dereferenced and a kernel panic occurred. This update adds the `flat_probe()` function to the APIC driver, which allows the kernel using flat APIC mode as a fall-back option. The kernel no longer panics in this situation.

BZ#952785

When attempting to deploy a virtual machine on a hypervisor with multiple NICs and macvtap devices, a kernel panic could occur. This happened because the macvtap driver did not gracefully handle a situation when the `macvlan_port.vlans` list was empty and returned a NULL pointer. This update applies a series of patches which fix this problem using a read-copy-update (RCU) mechanism and by preventing the driver from returning a NULL pointer if the list is empty. The kernel no longer panics in this scenario.

BZ#952329

Due to a missing structure, the NFSv4 error handler did not handle exceptions caused by revoking NFSv4 delegations. Consequently, the NFSv4 client received the EIO error message instead of the NFS4ERR_ADMIN_REVOKED error. This update modifies the NFSv4 code to no longer require the `nfs4_state` structure in order to revoke a delegation.

BZ#952174

On KVM guests with the KVM clock (`kvmclock`) as a clock source and with some VCPUs pinned, certain VCPUs could experience significant sleep delays (elapsed time was greater 20 seconds). This resulted in unexpected delays by sleeping functions and inaccurate measurement for low latency events. The problem happened because a `kvmclock` update was isolated to a certain VCPU so the NTP frequency correction applied only to that single VCPU. This problem has been resolved by a patch allowing `kvmclock` updates to all VCPUs on the KVM guest. VCPU sleep time now does not exceed the expected amount and no longer causes the aforementioned problems.

BZ#951937

When using applications that intensively utilized memory mapping, customers experienced significant application latency, which led to serious performance degradation. A series of patches has been applied to fix the problem. Among other, the patches modifies the memory mapping code to allow block devices to require stable page writes, enforce stable page writes only if required by a backing device, and optionally snapshot page content to provide stable pages during write. As a result, application latency has been improved by a considerable amount and applications with high demand of memory mapping now perform as expected.

BZ#997845

The RAID1 and RAID10 code previously called the `raise_barrier()` and `lower_barrier()` functions instead of the `freeze_array()` and `unfreeze_array()` functions that are safe being called from within the management thread. As a consequence, a deadlock situation could occur if an MD array contained a spare disk, rendering the respective kernel thread unresponsive. Furthermore, if a shutdown sequence was initiated after this problem had occurred, the shutdown sequence became unresponsive and any in-cache file system data that were not synchronized to the disk were lost. A patch correcting this problem has been applied and the RAID1 and RAID10 code now uses management-thread safe functions as expected.

BZ#996802

Previous changes to the Linux kernel network driver code introduced the TCP Small Queues (TSQ) feature. However, these changes led to performance degradation on certain network devices, such as devices using the `ixgbe` driver. This problem has been fixed by a series of patches to the TCP

Segmentation Offload (TSO) and TSQ features that include support for setting the size of TSO frames, and dynamic limit for the number of packet queues on device queues for a given TCP flow.

BZ#950598

If an NFSv4 client was checking open permissions for a delegated OPEN operation during OPEN state recovery of an NFSv4 server, the NFSv4 state manager could enter a deadlock. This happened because the client was holding the NFSv4 sequence ID of the OPEN operation. This problem is resolved by releasing the sequence ID before the client starts checking open permissions.

BZ#983288

NFS previously allowed extending an NFS file write to cover a full page only if the file had not set a byte-range lock. However, extending the write to cover the entire page is sometimes desirable in order to avoid fragmentation inefficiencies. For example, a noticeable performance decrease was reported if a series of small non-contiguous writes was performed on the file. A patch has been applied to the NFS code that allows NFS extending a file write to a full page write if the whole file is locked for writing or if the client holds a write delegation.

BZ#998752

A patch included in kernel version 2.6.32-358.9.1.el6, to fix handling of revoked NFSv4 delegations, introduced a regression bug to the NFSv4 code. This regression in the NFSv4 exception and asynchronous error handling allowed, under certain circumstances, passing a NULL inode to an NFSv4 delegation-related function, which resulted in a kernel panic. The NFSv4 exception and asynchronous error handling has been fixed so that a NULL inode can no longer be passed in this situation.

BZ#947582

XFS file systems were occasionally shut down with the "xfs_trans_ail_delete_bulk: attempting to delete a log item that is not in the AIL" error message. This happened because the EFI/EFD handling logic was incorrect and the EFI log item could have been freed before it was placed in the AIL and committed. A patch has been applied to the XFS code fixing the EFI/EFD handling logic and ensuring that the EFI log items are never freed before the EFD log items are processed. The aforementioned error no longer occurs on an XFS shutdown.

BZ#947275

A bug in the autofs4 mount expiration code could cause the autofs4 module to falsely report a busy tree of NFS mounts as "not in use". Consequently, automount attempted to unmount the tree and failed with a "failed to umount offset" error, leaving the mount tree to appear as empty directories. A patch has been applied to remove an incorrectly used autofs dentry mount check and the aforementioned problem no longer occurs.

BZ#927988

Cyclic adding and removing of the st kernel module could previously cause a system to become unresponsive. This was caused by a disk queue reference count bug in the SCSI tape driver. An upstream patch addressing this bug has been backported to the SCSI tape driver and the system now responds as expected in this situation.

BZ#927918

A previous update introduced a new failure mode to the blk_get_request() function returning the -ENODEV error code when a block device queue is being destroyed. However, the change did not include a NULL pointer check for all callers of the function. Consequently, the kernel could dereference a NULL pointer when removing a block device from the system, which resulted in a kernel panic. This update applies a patch that adds these missing NULL pointer checks. Also, some

callers of the `blk_get_request()` function could previously return the `-ENOMEM` error code instead of `-ENODEV`, which would lead to incorrect call chain propagation. This update applies a patch ensuring that correct return codes are propagated.

BZ#790921

By default, the kernel uses a best-fit algorithm for allocating Virtual Memory Areas (VMAs) to map processed files to the address space. However, if an enormous number of small files (hundreds of thousands or millions) was being mapped, the address space became extremely fragmented, which resulted in significant CPU usage and performance degradation. This update introduces an optional next-fit policy which, if enabled, allows for mapping of a file to the first suitable unused area in the address space that follows after the previously allocated VMA.

BZ#960717

A rare race condition between the "devloss" timeout and discovery state machine could trigger a bug in the `lpfc` driver that nested two levels of spin locks in reverse order. The reverse order of spin locks led to a deadlock situation and the system became unresponsive. With this update, a patch addressing the deadlock problem has been applied and the system no longer hangs in this situation.

BZ#922999

An error in backporting the block reservation feature from upstream resulted in a missing allocation of a reservation structure when an allocation is required during the rename system call. Renaming a file system object (for example, file or directory) requires a block allocation for the destination directory. If the destination directory had not had a reservation structure allocated, a NULL pointer dereference occurred, leading to a kernel panic. With this update, a reservation structure is allocated before the rename operation, and a kernel panic no longer occurs in this scenario.

BZ#805407

A system could become unresponsive due to an attempt to shut down an XFS file system that was waiting for log I/O completion. A patch to the XFS code has been applied that allows for the shutdown method to be called from different contexts so XFS log items can be deleted properly even outside the AIL, which fixes this problem.

BZ#922931

A bug in the `dm_btree_remove()` function could cause leaf values to have incorrect reference counts. Removal of a shared block could result in space maps considering the block as no longer used. As a consequence, sending a discard request to a shared region of a thin device could corrupt its snapshot. The bug has been fixed to prevent corruption in this scenario.

BZ#980273

A recent change in the memory mapping code introduced a new optional next-fit algorithm for allocating VMAs to map processed files to the address space. This change, however, broke behavior of a certain internal function which then always followed the next-fit VMA allocation scheme instead of the first-fit VMA allocation scheme. Consequently, when the first-fit VMA allocation scheme was in use, this bug caused linear address space fragmentation and could lead to early `"-ENOMEM"` failures for `mmap()` requests. This patch restores the original first-fit behavior to the function so the aforementioned problems no longer occur.

BZ#922779

The GFS2 discard code did not calculate the sector offset correctly for block devices with the sector size of 4 KB, which led to loss of data and metadata on these devices. A patch correcting this problem has been applied so the discard and FITRIM requests now work as expected for the block devices with the 4 KB sector size.

BZ#1002765

A bug in the real-time (RT) scheduler could cause a RT priority process to stop running due to an invalid attribute of the run queue. When a CPU became affected by this bug, the migration kernel thread stopped running on the CPU, and subsequently every other process that was migrated to the affected CPU by the system stopped running as well. A patch has been applied to the RT scheduler and RT priority processes are no longer affected this problem.

BZ#920794

When using the congestion window lock functionality of the ip utility, the system could become unresponsive. This happened because the `tcp_slow_start()` function could enter an infinite loop if the congestion window was locked using route metrics. A set of patches has been applied to comply with the upstream kernel, ensuring the problem no longer occurs in this scenario.

BZ#978609

A race condition in the abort task and SPP device task management path of the iscsi driver could, under certain circumstances, cause the driver to fail cleaning up timed-out I/O requests that were pending on an SAS disk device. As a consequence, the kernel removed such a device from the system. A patch applied to the iscsi driver fixes this problem by sending the task management function request to the SAS drive anytime the abort function is entered and the task has not completed. The driver now cleans up timed-out I/O requests as expected in this situation.

BZ#920672

Due to a race condition in the kernel's DMA initialization code, DMA requests from the `hpsa` and `hpilo` drivers could fail with `IO_PAGE_FAULT` errors during initialization of the AMD iommu driver on AMD systems with the IOMMU feature enabled. To avoid triggering this race condition, the kernel now executes the `init_device_table_dma()` function to block DMA requests from all devices only after the initialization of unity mappings is finished.

BZ#1003697

If the `arp_interval` and `arp_validate` bonding options were not enabled on the configured bond device in the correct order, the bond device did not process ARP replies, which led to link failures and changes of the active slave device. A series of patches has been applied to modify an internal bond ARP hook based on the values of `arp_validate` and `arp_interval`. Therefore, the ARP hook is registered even if `arp_interval` is set after `arp_validate` has already been enabled, and ARP replies are processed as expected.

BZ#920445

The kernel could rarely terminate instead of creating a dump file when a multi-threaded process using FPU aborted. This happened because the kernel did not wait until all threads became inactive and attempted to dump the FPU state of active threads into memory which triggered a `BUG_ON()` routine. A patch addressing this problem has been applied and the kernel now waits for the threads to become inactive before dumping their FPU state into memory.

BZ#962460

Previously, the Generic Receive Offload (GRO) functionality was not enabled by default for VLAN devices. Consequently, certain network adapters, such as Emulex Virtual Fabric Adapter (VFA) II, that use `be2net` driver, were dropping packets when VLAN tagging was enabled and the `8021q` kernel module loaded. This update applies a patch that enables GRO by default for VLAN devices.

BZ#827548

A race condition between the `read_swap_cache_async()` and `get_swap_page()` functions in the

Memory management (mm) code could lead to a deadlock situation. The deadlock could occur only on systems that deployed swap partitions on devices supporting block DISCARD and TRIM operations if kernel preemption was disabled (the `!CONFIG_PREEMPT` parameter). If the `read_swap_cache_async()` function was given a `SWAP_HAS_CACHE` entry that did not have a page in the swap cache yet, a DISCARD operation was performed in the `scan_swap_map()` function. Consequently, completion of an I/O operation was scheduled on the same CPU's working queue the `read_swap_cache_async()` was running on. This caused the thread in `read_swap_cache_async()` to loop indefinitely around its `"-EEXIST"` case, rendering the system unresponsive. The problem has been fixed by adding an explicit `cond_resched()` call to `read_swap_cache_async()`, which allows other tasks to run on the affected CPU, and thus avoiding the deadlock.

BZ#987426

An infinite loop bug in the NFSv4 code caused an NFSv4 mount process to hang on a busy loop of the `LOOKUP_ROOT` operation when attempting to mount an NFSv4 file system and the first iteration on this operation failed. A patch has been applied that allows to exit the `LOOKUP_ROOT` operation properly and a mount attempt now either succeeds or fails in this situation.

BZ#828936

A bug in the OProfile tool led to a NULL pointer dereference while unloading the OProfile kernel module, which resulted in a kernel panic. The problem was triggered if the kernel was running with the `nolapic` parameter set and OProfile was configured to use the NMI timer interrupt. The problem has been fixed by correctly setting the NMI timer when initializing OProfile.

BZ#976915

An NFS client previously did not wait for completing of unfinished I/O operations before sending the `LOCKU` and `RELEASE_LOCKOWNER` operations to the NFS server in order to release byte range locks on files. Consequently, if the server processed the `LOCKU` and `RELEASE_LOCKOWNER` operations before some of the related `READ` operations, it released all locking states associated with the requested lock owner, and the `READs` returned the `NFS4ERR_BAD_STATEID` error code. This resulted in the "Lock reclaim failed!" error messages being generated in the system log and the NFS client had to recover from the error. A series of patches has been applied ensuring that an NFS client waits for all outstanding I/O operations to complete before releasing the locks.

BZ#918239

When the Red Hat Enterprise Linux 6 kernel runs as a virtual machine, it performs boot-time detection of the hypervisor in order to enable hypervisor-specific optimizations. Red Hat Enterprise Linux 6.4 introduces detection and optimization for the Microsoft Hyper-V hypervisor. Previously Hyper-V was detected first, however, because some Xen hypervisors can attempt to emulate Hyper-V, this could lead to a boot failure when that emulation was not exact. A patch has been applied to ensure that the attempt to detect Xen is always done before Hyper-V, resolving this issue.

BZ#962976

If the audit queue is too long, the kernel schedules the `kauditd` daemon to alleviate the load on the audit queue. Previously, if the current audit process had any pending signals in such a situation, it entered a busy-wait loop for the duration of an audit backlog timeout because the `wait_for_auditd()` function was called as an interruptible task. This could lead to system lockup in non-preemptive uniprocessor systems. This update fixes the problem by setting `wait_for_auditd()` as uninterruptible.

BZ#833299

Due to a bug in firmware, systems using the LSI MegaRAID controller failed to initialize this device in the `kdump` kernel if the `"intel_iommu=on"` and `"iommu=pt"` kernel parameters were specified in the first kernel. As a workaround until a firmware fix is available, a patch to the `megaraid_sas` driver has been

applied so if the firmware is not in the ready state upon the first attempt to initialize the controller, the driver resets the controller and retries for firmware transition to the ready state.

BZ#917872

A previous change in the port auto-selection code allowed sharing ports with no conflicts extending its usage. Consequently, when binding a socket with the `SO_REUSEADDR` socket option enabled, the `bind(2)` function could allocate an ephemeral port that was already used. A subsequent connection attempt failed in such a case with the `EADDRNOTAVAIL` error code. This update applies a patch that modifies the port auto-selection code so that `bind(2)` now selects a non-conflict port even with the `SO_REUSEADDR` option enabled.

BZ#994430

A previous patch to the bridge multicast code introduced a bug allowing reinitialization of an active timer for a multicast group whenever an IPv6 multicast query was received. A patch has been applied to the bridge multicast code so that a bridge multicast timer is no longer reinitialized when it is active.

BZ#916994

A kernel panic could occur during path failover on systems using multiple iSCSI, FC or SRP paths to connect an iSCSI initiator and an iSCSI target. This happened because a race condition in the SCSI driver allowed removing a SCSI device from the system before processing its run queue, which led to a NULL pointer dereference. The SCSI driver has been modified and the race is now avoided by holding a reference to a SCSI device run queue while it is active.

BZ#994382

The kernel's md driver contained multiple bugs, including a use-after-free bug in the raid10 code that could cause a kernel panic. Also a data corruption bug in the raid5 code was discovered. The bug occurred when a hard drive was replaced while a RAID4, RAID5, or RAID6 array contained by the drive was in process of recovery. A series of patches has been applied to fix all bugs that have been discovered. The md driver now contains necessary tests that prevent the mentioned use-after-free and data corruption bugs from occurring.

BZ#840860

The `sunrpc` code paths that wake up an RPC task are highly optimized for speed so the code avoids using any locking mechanism but requires precise operation ordering. Multiple bugs were found related to operation ordering, which resulted in a kernel crash involving either a `BUG_ON()` assertion or an incorrect use of a data structure in the `sunrpc` layer. These problems have been fixed by properly ordering operations related to the `RPC_TASK_QUEUED` and `RPC_TASK_RUNNING` bits in the wake-up code paths of the `sunrpc` layer.

BZ#916735

In the RPC code, when a network socket backed up due to high network traffic, a timer was set causing a retransmission, which in turn could cause even larger amount of network traffic to be generated. To prevent this problem, the RPC code now waits for the socket to empty instead of setting the timer.

BZ#916726

When using parallel NFS (pNFS), a kernel panic could occur when a process was killed while getting the file layout information during the `open()` system call. A patch has been applied to prevent this problem from occurring in this scenario.

BZ#916722

Previously, when `open(2)` system calls were processed, the `GETATTR` routine did not check to see if valid attributes were also returned. As a result, the `open()` call succeeded with invalid attributes instead of failing in such a case. This update adds the missing check, and the `open()` call succeeds only when valid attributes are returned.

BZ#916361

The `crypto_larval_lookup()` function could return a larval, an in-between state when a cryptographic algorithm is being registered, even if it did not create one. This could cause a larval to be terminated twice, and result in a kernel panic. This occurred for example when the NFS service was run in FIPS mode, and attempted to use the MD5 hashing algorithm even though FIPS mode has this algorithm blacklisted. A condition has been added to the `crypto_larval_lookup()` function to check whether a larval was created before returning it.

BZ#976879

Previously, systems running heavily-loaded NFS servers could experience poor performance of the NFS `REaddir` operations on large directories that were undergoing concurrent modifications, especially over higher latency connections. This happened because the NFS code performed certain dentry operations inefficiently and revalidated directory attributes too often. This update applies a series of patches that address the problem as follows; needed dentries can be accessed from `dcache` after the `REaddir` operation, and directory attributes are revalidated only at the beginning of the directory or if the cached attributes expire.

BZ#976823

The GFS2 did not reserve journal space for a quota change block while growing the size of a file. Consequently, a fatal assertion causing a withdraw of the GFS2 file system could have been triggered when the free blocks were allocated from the secondary bitmap. With this update, GFS2 reserves additional blocks in the journal for the quota change so the file growing transaction can now complete successfully in this situation.

BZ#976535

A previous patch to the CIFS code caused a regression of a problem where under certain conditions, a mount attempt of a CIFS DFS share fails with a "mount error(6): No such device or address" error message. This happened because the return code variable was not properly reset after a previous unsuccessful mount attempt. A backported patch has been applied to properly reset the variable and CIFS DFS shares can now be mounted as expected.

BZ#965002

A bug in the PCI driver allowed to use a pointer to the Virtual Function (VF) device entry that was already freed. Consequently, when hot-removing an I/O unit with enabled SR-IOV devices, a kernel panic occurred. This update modifies the PCI driver so a valid pointer to the Physical Function (PF) device entry is used and the kernel no longer panics in this situation.

BZ#915834

A race condition could occur in the `uhci-hcd` kernel module if the IRQ line was shared with other devices. The race condition allowed the IRQ handler routine to be called before the data structures were fully initialized, which caused the system to become unresponsive. This update applies a patch that fixes the problem by adding a test condition to the IRQ handler routine; if the data structure initialization is still in progress, the handler routine finishes immediately.

BZ#975507

An insufficiently designed calculation in the CPU accelerator could cause an arithmetic overflow in the `set_cyc2ns_scale()` function if the system uptime exceeded 208 days prior to using `kexec` to boot into

a new kernel. This overflow led to a kernel panic on the systems using the Time Stamp Counter (TSC) clock source, primarily the systems using Intel Xeon E5 processors that do not reset TSC on soft power cycles. A patch has been applied to modify the calculation so that this arithmetic overflow and kernel panic can no longer occur under these circumstances.

BZ#915479

Due to a bug in the NFSv4 `nfsd` code, a NULL pointer could have been dereferenced when `nfsd` was looking up a path to the NFSv4 recovery directory for the `fsync` operation, which resulted in a kernel panic. This update applies a patch that modifies the NFSv4 `nfsd` code to open a file descriptor for `fsync` in the NFSv4 recovery directory instead of looking up the path. The kernel no longer panics in this situation.

BZ#858198

Previously, bond and bridge devices did not pass Generic Receive Offload (GRO) information to their slave devices, and bridge devices also did not propagate VLAN information to their ports. As a consequence, in environments with VLAN configured over a bridge or bonding device, performance of the slave devices configured on the bridge and bonding devices was significantly low. A series of patches has been applied that adds the GRO feature for bonding and bridge devices and allows VLANs to be registered with the participating bridge ports. If a slave device supports GRO, its performance is now significantly increased in environments with VLAN configured over a bridge or bonding device.

BZ#975211

Due to a bug in the NFS code, kernel size-192 and size-256 slab caches could leak memory. This could eventually result in an OOM issue when the most of available memory was used by the respective slab cache. A patch has been applied to fix this problem and the respective attributes in the NFS code are now freed properly.

BZ#913704

Previously, the NFS Lock Manager (NLM) did not resend blocking lock requests after NFSv3 server reboot recovery. As a consequence, when an application was running on a NFSv3 mount and requested a blocking lock, the application received an `-ENOLCK` error. This patch ensures that NLM always resend blocking lock requests after the grace period has expired.

BZ#862758

When counting CPU time, the `utime` and `stime` values are scaled based on `rtime`. Prior to this update, the `utime` value was multiplied with the `rtime` value, but the integer multiplication overflow could happen, and the resulting value could be then truncated to 64 bits. As a consequence, `utime` values visible in the user space were stall even if an application consumed a lot of CPU time. With this update, the multiplication is performed on `stime` instead of `utime`. This significantly reduces the chances of an overflow on most workloads because the `stime` value, unlike the `utime` value, cannot grow fast.

BZ#913660

In a case of a broken or malicious server, an index node (`inode`) of an incorrect type could be matched. This led to an NFS client NULL pointer dereference, and, consequently, to a kernel oops. To prevent this problem from occurring in this scenario, a check has been added to verify that the `inode` type is correct.

BZ#913645

A previously-applied patch introduced a bug in the `ipoib_cm_destroy_tx()` function, which allowed a CM object to be moved between lists without any supported locking. Under a heavy system load, this

could cause the system to crash. With this update, proper locking of the CM objects has been re-introduced to fix the race condition, and the system no longer crashes under a heavy load.

BZ#966853

Previously, when booting a Red Hat Enterprise Linux 6.4 system and the ACPI Static Resource Affinity Table (SRAT) had a hot-pluggable bit enabled, the kernel considered the SRAT table incorrect and NUMA was not configured. This led to a general protection fault and a kernel panic occurring on the system. The problem has been fixed by using an SMBIOS check in the code in order to avoid the SRAT code table consistency checks. NUMA is now configured as expected and the kernel no longer panics in this situation.

BZ#912963

When booting the normal kernel on certain servers, such as HP ProLiant DL980 G7, some interrupts may have been lost which resulted in the system being unresponsive or rarely even in data loss. This happened because the kernel did not set correct destination mode during the boot; the kernel booted in "logical cluster mode" that is default while this system supported only "x2apic physical mode". This update applies a series of patches addressing the problem. The underlying APIC code has been modified so the x2apic probing code now checks the Fixed ACPI Description Table (FADT) and installs the x2apic "physical" driver as expected. Also, the APIC code has been simplified and the code now uses probe routines to select destination APIC mode and install the correct APIC drivers.

BZ#912867

Previously, the `fsync(2)` system call incorrectly returned the EIO (Input/Output) error instead of the ENOSPC (No space left on device) error. This was due to incorrect error handling in the page cache. This problem has been fixed and the correct error value is now returned.

BZ#912842

Previously, an NFS RPC task could enter a deadlock and become unresponsive if it was waiting for an NFSv4 state serialization lock to become available and the session slot was held by the NFSv4 server. This update fixes this problem along with the possible race condition in the pNFS return-on-close code. The NFSv4 client has also been modified to not accepting delegated OPEN operations if a delegation recall is in effect. The client now also reports NFSv4 servers that try to return a delegation when the client is using the CLAIM_DELEGATE_CUR open mode.

BZ#912662

Due to the way the CPU time was calculated, an integer multiplication overflow bug could occur after several days of running CPU bound processes that were using hundreds of kernel threads. As a consequence, the kernel stopped updating the CPU time and provided an incorrect CPU time instead. This could confuse users and lead to various application problems. This update applies a patch fixing this problem by decreasing the precision of calculations when the `stime` and `rtime` values become too large. Also, a bug allowing `stime` values to be sometimes erroneously calculated as `utime` values has been fixed.

BZ#967095

An NFS server could terminate unexpectedly due to a NULL pointer dereference caused by a rare race condition in the `lockd` daemon. An applied patch fixes this problem by protecting the relevant code with spin locks, and thus avoiding the race in `lockd`.

BZ#911359

Virtual LAN (VLAN) support of the eHEA ethernet adapter did not work as expected. A "device ethX has buggy VLAN hw accel" message could have been reported when running the "dmesg" command. This was because an upstream backport patch removed the `vlan_rx_register()` function. This update

adds the function back, and eHEA VLAN support works as expected. This update also addresses a possible kernel panic, which could occur due to a NULL pointer dereference when processing received VLAN packets. The patch adds a test condition verifying whether a VLAN group is set by the network stack, which prevents a possible NULL pointer to be dereferenced, and the kernel no longer crashes in this situation.

BZ#910597

The kernel's implementation of RTAS (RunTime Abstraction Services) previously allowed the `stop_topology_update()` function to be called from an interrupt context during live partition migration on PowerPC and IBM System p machines. As a consequence, the system became unresponsive. This update fixes the problem by calling `stop_topology_update()` earlier in the migration process, and the system no longer hangs in this situation.

BZ#875753

Truncating files on a GFS2 file system could fail with an "unable to handle kernel NULL pointer dereference" error. This was because of a missing reservation structure that caused the truncate code to reference an incorrect pointer. To prevent this, a patch has been applied to allocate a block reservation structure before truncating a file.

BZ#909464

Previously, race conditions could sometimes occur in interrupt handling on the Emulex BladeEngine 2 (BE2) controllers, causing the network adapter to become unresponsive. This update provides a series of patches for the `be2net` driver, which prevents the race from occurring. The network cards using BE2 chipsets no longer hang due to incorrectly handled interrupt events.

BZ#908990

Previously, power-limit notification interrupts were enabled by default on the system. This could lead to degradation of system performance or even render the system unusable on certain platforms, such as Dell PowerEdge servers. A patch has been applied to disable power-limit notification interrupts by default and a new kernel command line parameter "`int_pln_enable`" has been added to allow users observing these events using the existing system counters. Power-limit notification messages are also no longer displayed on the console. The affected platforms no longer suffer from degraded system performance due to this problem.

BZ#876778

A change in the `ipmi_si` driver handling caused an extensively long delay while booting Red Hat Enterprise Linux 6.4 on SIG UV platforms. The driver was loaded as a kernel module on previous versions of Red Hat Enterprise Linux 6 while it is now built within the kernel. However, SIG UV does not use, and thus does not support the `ipmi_si` driver. A patch has been applied and the kernel now does not initialize the `ipmi_si` driver when booting on SIG UV.

BZ#908851

Previously, the queue limits were not being retained as they should have been if a device did not contain any data or if a multipath device temporarily lost all its paths. This problem has been fixed by avoiding a call to the `dm_calculate_queue_limits()` function.

BZ#908751

When adding a virtual PCI device, such as `virtio disk`, `virtio net`, `e1000` or `rtl8139`, to a KVM guest, the `kacpid` thread reprograms the hot plug parameters of all devices on the PCI bus to which the new device is being added. When reprogramming the hot plug parameters of a VGA or QXL graphics device, the graphics device emulation requests flushing of the guest's shadow page tables. Previously, if the guest had a huge and complex set of shadow page tables, the flushing operation

took a significant amount of time and the guest could appear to be unresponsive for several minutes. This resulted in exceeding the threshold of the "soft lockup" watchdog and the "BUG: soft lockup" events were logged by both, the guest and host kernel. This update applies a series of patches that deal with this problem. The KVM's Memory Management Unit (MMU) now avoids creating multiple page table roots in connection with processors that support Extended Page Tables (EPT). This prevents the guest's shadow page tables from becoming too complex on machines with EPT support. MMU now also flushes only large memory mappings, which alleviates the situation on machines where the processor does not support EPT. Additionally, a free memory accounting race that could prevent KVM MMU from freeing memory pages has been fixed.

BZ#908608

Certain CPUs contain on-chip virtual-machine control structure (VMCS) caches that are used to keep active VMCSs managed by the KVM module. These VMCSs contain runtime information of the guest machines operated by KVM. These CPUs require support of the VMCLEAR instruction that allows flushing the cache's content into memory. The kernel previously did not use the VMCLEAR instruction in Kdump. As a consequence, when dumping a core of the QEMU KVM host, the respective CPUs did not flush VMCSs to the memory and the guests' runtime information was not included in the core dump. This problem has been addressed by a series of patches that implement support of using the VMCLEAR instruction in Kdump. The kernel is now performs the VMCLEAR operation in Kdump if it is required by a CPU so the vmcore file of the QEMU KVM host contains all VMCSs information as expected.

BZ#908524

When pNFS (parallel NFS) code was in use, a file locking process could enter a deadlock while trying to recover from a server reboot. This update introduces a new locking mechanism that avoids the deadlock situation in this scenario.

BZ#878708

Sometimes, the irqbalance tool could not get the CPU NUMA node information because of missing symlinks for CPU devices in sysfs. This update adds the NUMA node symlinks for CPU devices in sysfs, which is also useful when using irqbalance to build a CPU topology.

BZ#908158

The virtual file system (VFS) code had a race condition between the unlink and link system calls that allowed creating hard links to deleted (unlinked) files. This could, under certain circumstances, cause inode corruption that eventually resulted in a file system shutdown. The problem was observed in Red Hat Storage during rsync operations on replicated Gluster volumes that resulted in an XFS shutdown. A testing condition has been added to the VFS code, preventing hard links to deleted files from being created.

BZ#908093

When an inconsistency is detected in a GFS2 file system after an I/O operation, the kernel performs the withdraw operation on the local node. However, the kernel previously did not wait for an acknowledgement from the GFS control daemon (`gfs_controld`) before proceeding with the withdraw operation. Therefore, if a failure isolating the GFS2 file system from a data storage occurred, the kernel was not aware of this problem and an I/O operation to the shared block device may have been performed after the withdraw operation was logged as successful. This could lead to corruption of the file system or prevent the node from journal recovery. This patch modifies the GFS2 code so the withdraw operation no longer proceeds without the acknowledgement from `gfs_controld`, and the GFS2 file system can no longer become corrupted after performing the withdraw operation.

BZ#907844

If a logical volume was created on devices with thin provisioning enabled, the `mkfs.ext4` command took a long time to complete, and the following message was recorded in the system log:

```
kernel: blk: request botched
```

This was caused by discard request merging that was not completely functional in the block and SCSI layers. This functionality has been temporarily disabled to prevent such problems from occurring.

BZ#907512

A previous patch that modified `dcache` and `autofs` code caused a regression. Due to this regression, unmounting a large number of expired automounts on a system under heavy NFS load caused soft lockups, rendering the system unresponsive. If a "soft lockup" watchdog was configured, the machine rebooted. To fix the regression, the erroneous patch has been reverted and the system now handle the aforementioned scenario properly without any soft lockups.

BZ#907227

Previously, when using parallel network file system (pNFS) and data was written to the appropriate storage device, the `LAYOUTCOMMIT` requests being sent to the metadata server could fail internally. The metadata server was not provided with the modified layout based on the written data, and these changes were not visible to the NFS client. This happened because the encoding functions for the `LAYOUTCOMMIT` and `LAYOUTRETURN` operations were defined as void, and returned thus an arbitrary status. This update corrects these encoding functions to return 0 on success as expected. The changes on the storage device are now propagated to the metadata server and can be observed as expected.

BZ#883905

When the Active Item List (AIL) becomes empty, the `xfsaild` daemon is moved to a task sleep state that depends on the timeout value returned by the `xfsaild_push()` function. The latest changes modified `xfsaild_push()` to return a 10-ms value when the AIL is empty, which sets `xfsaild` into the uninterruptible sleep state (D state) and artificially increased system load average. This update applies a patch that fixes this problem by setting the timeout value to the allowed maximum, 50 ms. This moves `xfsaild` to the interruptible sleep state (S state), avoiding the impact on load average.

BZ#905126

Previously, init scripts were unable to set the master interface MAC address properly because it was overwritten by the first slave MAC address. To avoid this problem, this update re-introduces the check for an unassigned MAC address before adopting the first slaves as its own.

BZ#884442

Due to a bug in the `be2net` driver, events in the RX, TX, and MCC queues were not acknowledged before closing the respective queue. This could cause unpredictable behavior when creating RX rings during the subsequent queue opening. This update applies a patch that corrects this problem and events are now acknowledged as expected in this scenario.

BZ#904726

Previously, the `mlx4` driver set the number of requested MSI-X vectors to 2 under multi-function mode on `mlx4` cards. However, the default setting of the `mlx4` firmware allows for a higher number of requested MSI-X vectors (4 of them with the current firmware). This update modifies the `mlx4` driver so that it uses these default firmware settings, which improves performance of `mlx4` cards.

BZ#904025

Reading a large number of files from a pNFS (parallel NFS) mount and canceling the running operation by pressing Ctrl+C caused a general protection fault in the XDR code, which could manifest itself as a kernel oops with an "unable to handle kernel paging request" message. This happened because decoding of the LAYOUTGET operation is done by a worker thread and the caller waits for the worker thread to complete. When the reading operation was canceled, the caller stopped waiting and freed the pages. So the pages no longer existed at the time the worker thread called the relevant function in the XDR code. The cleanup process of these pages has been moved to a different place in the code, which prevents the kernel oops from happening in this scenario.

BZ#903644

A previous patch to the mlx4 driver enabled an internal loopback to allow communication between functions on the same host. However, this change introduced a regression that caused virtual switch (vSwitch) bridge devices using Mellanox Ethernet adapter as the uplink to become inoperative in native (non-SRIOV) mode under certain circumstances. To fix this problem, the destination MAC address is written to Tx descriptors of transmitted packets only in SRIOV or eSwitch mode, or during the device self-test. Uplink traffic works as expected in the described setup.

BZ#887006

The Intel 5520 and 5500 chipsets do not properly handle remapping of MSI and MSI-X interrupts. If the interrupt remapping feature is enabled on the system with such a chipset, various problems and service disruption could occur (for example, a NIC could stop receiving frames), and the "kernel: do_IRQ: 7.71 No irq handler for vector (irq -1)" error message appears in the system logs. As a workaround to this problem, it has been recommended to disable the interrupt remapping feature in the BIOS on such systems, and many vendors have updated their BIOS to disable interrupt remapping by default. However, the problem is still being reported by users without proper BIOS level with this feature properly turned off. Therefore, this update modifies the kernel to check if the interrupt remapping feature is enabled on these systems and to provide users with a warning message advising them on turning off the feature and updating the BIOS.

BZ#887045

When booting Red Hat Enterprise Linux 6 system that utilized a large number of CPUs (more than 512), the system could fail to boot or could appear to be unresponsive after initialization. This happened because the CPU frequency driver used a regular spin lock (cpufreq_driver_lock) to serialize frequency transitions, and this lock could, under certain circumstances, become a source of heavy contention during the system initialization and operation. A patch has been applied to convert cpufreq_driver_lock into a read-write lock, which resolves the contention problem. All Red Hat Enterprise Linux 6 systems now boot and operate as expected.

BZ#903220

A previous patch to the kernel introduced a bug by assigning a different value to the IFLA_EXT_MASK Netlink attribute than found in the upstream kernels. This could have caused various problems; for example, a binary compiled against upstream headers could have failed or behaved unexpectedly on Red Hat Enterprise Linux 6.4 and later kernels. This update realigns IFLA_EXT_MASK in the enumeration correctly by synchronizing the IFLA_* enumeration with the upstream. This ensures that binaries compiled against Red Hat Enterprise Linux 6.4 kernel headers will function as expected. Backwards compatibility is guaranteed.

BZ#887868

Due to a bug in the SCTP code, a NULL pointer dereference could occur when freeing an SCTP association that was hashed, resulting in a kernel panic. A patch addresses this problem by trying to unhash SCTP associations before freeing them and the problem no longer occurs.

BZ#888417

Previously, a kernel panic could occur on machines using the SCSI sd driver with Data Integrity Field (DIF) type 2 protection. This was because the `scsi_register_driver()` function registered the `prep_fn()` function that might have needed to use the `sd_cdp_pool` variable for the DIF functionality. However, the variable had not yet been initialized at this point. The underlying code has been updated so that the driver is registered last, which prevents a kernel panic from occurring in this scenario.

BZ#901747

The `bnx2x` driver could have previously reported an occasional MDC/MDIO timeout error along with the loss of the link connection. This could happen in environments using an older boot code because the MDIO clock was set in the beginning of each boot code sequence instead of per CL45 command. To avoid this problem, the `bnx2x` driver now sets the MDIO clock per CL45 command. Additionally, the MDIO clock is now implemented per EMAC register instead of per port number, which prevents ports from using different EMAC addresses for different PHY accesses. Also, boot code or Management Firmware (MFW) upgrade is required to prevent the boot code (firmware) from taking over link ownership if the driver's pulse is delayed. The BCM57711 card requires boot code version 6.2.24 or later, and the BCM57712/578xx cards require MFW version 7.4.22 or later.

BZ#990806

When the Audit subsystem was under heavy load, it could loop infinitely in the `audit_log_start()` function instead of failing over to the error recovery code. This would cause soft lockups in the kernel. With this update, the timeout condition in the `audit_log_start()` function has been modified to properly fail over when necessary.

BZ#901701

A previous kernel update broke queue pair (qp) hash list deletion in the `qp_remove()` function. This could cause a general protection fault in the InfiniBand stack or QLogic InfiniBand driver. A patch has been applied to restore the former behavior so the general protection fault no longer occurs.

BZ#896233

Under rare circumstances, if a TCP retransmission was multiple times partially acknowledged and collapsed, the used Socked Buffer (SKB) could become corrupted due to an overflow caused by the transmission headroom. This resulted in a kernel panic. The problem was observed rarely when using an IP-over-InfiniBand (IPoIB) connection. This update applies a patch that verifies whether a transmission headroom exceeded the maximum size of the used SKB, and if so, the headroom is reallocated. It was also discovered that a TCP stack could retransmit misaligned SKBs if a malicious peer acknowledged sub MSS frame and output interface did not have a sequence generator (SG) enabled. This update introduces a new function that allows for copying of a SKB with a new head so the SKB remains aligned in this situation.

BZ#896020

When using transparent proxy (TProxy) over IPv6, the kernel previously created neighbor entries for local interfaces and peers that were not reachable directly. This update corrects this problem and the kernel no longer creates invalid neighbor entries.

BZ#894683

A previous change in the port auto-selection code allowed sharing ports with no conflicts extending its usage. Consequently, when binding a socket with the `SO_REUSEADDR` socket option enabled, the `bind(2)` function could allocate an ephemeral port that was already used. A subsequent connection attempt failed in such a case with the `EADDRNOTAVAIL` error code. This update applies a patch that modifies the port auto-selection code so that `bind(2)` now selects a non-conflict port even with the `SO_REUSEADDR` option enabled.

BZ#893584

Timeouts could occur on an NFS client with heavy read workloads; for example when using `rsync` and `ldconfig`. Both client-side and server-side causes were found for the problem. On the client side, problems that could prevent the client reconnecting lost TCP connections have been fixed. On the server side, TCP memory pressure on the server forced the send buffer size to be lower than the size required to send a single Remote Procedure Call (RPC), which consequently caused the server to be unable to reply to the client. Code fixes are still being considered. To work around the problem, increase the minimum TCP buffer sizes, for example using:

```
echo "1048576 1048576 4194304" >/proc/sys/net/ipv4/tcp_wmem
```

BZ#895336

Broadcom 5719 NIC could previously sometimes drop received jumbo frame packets due to cyclic redundancy check (CRC) errors. This update modifies the `tg3` driver so that CRC errors no longer occur and Broadcom 5719 NICs process jumbo frame packets as expected.

BZ#896224

When running a high thread workload of small-sized files on an XFS file system, sometimes, the system could become unresponsive or a kernel panic could occur. This occurred because the `xfsaild` daemon had a subtle code path that led to lock recursion on the `xfsaild` lock when a buffer in the AIL was already locked and an attempt was made to force the log to unlock it. This patch removes the dangerous code path and queues the log force to be invoked from a safe locking context with respect to `xfsaild`. This patch also fixes the race condition between buffer locking and buffer pinned state that exposed the original problem by rechecking the state of the buffer after a lock failure. The system no longer hangs and kernel no longer panics in this scenario.

BZ#902965

The NFSv4.1 client could stop responding while recovering from a server reboot on an NFSv4.1 or pNFS mount with delegations disabled. This could happen due to insufficient locking in the NFS code and several related bugs in the NFS and RPC scheduler code which could trigger a deadlock situation. This update applies a series of patches which prevent possible deadlock situations from occurring. The NFSv4.1 client now recovers and continue with workload as expected in the described situation.

BZ#1010840

The default `sfc` driver on Red Hat Enterprise Linux 6 allowed toggling the Large Receive Offset (LRO) flag on and off on a network device regardless of whether LRO was supported by the device or not. Therefore, when the LRO flag was enabled on devices without LRO support, the action had no effect and could confuse users. A patch to the `sfc` driver has been applied so that the `sfc` driver properly validates whether LRO is supported by the device. If the device does not support LRO, `sfc` disables the LRO flag so that users can no longer toggle it for that device.

BZ#886867

During device discovery, the system creates a temporary SCSI device with the LUN ID 0 if the LUN 0 is not mapped on the system. Previously, this led to a NULL pointer dereference because inquiry data was not allocated for the temporary LUN 0 device, which resulted in a kernel panic. This update adds a NULL pointer test in the underlying SCSI code, and the kernel no longer panics in this scenario.

BZ#886420

When a network interface (NIC) is running in promiscuous (PROMISC) mode, the NIC may receive and process VLAN tagged frames even though no VLAN is attached to the NIC. However, some

network drivers, such as bnx2, igb, tg3, and e1000e did not handle processing of packets with VLAN tagged frames in PROMISC mode correctly if the frames had no VLAN group assigned. The drivers processed the packets with incorrect routines and various problems could occur; for example, a DHCPv6 server connected to a VLAN could assign an IPv6 address from the VLAN pool to a NIC with no VLAN interface. To handle the VLAN tagged frames without a VLAN group properly, the frames have to be processed by the VLAN code so the aforementioned drivers have been modified to restrain from performing a NULL value test of the packet's VLAN group field when the NIC is in PROMISC mode. This update also includes a patch fixing a bug where the bnx2x driver did not strip a VLAN header from the frame if no VLAN was configured on the NIC, and another patch that implements some register changes in order to enable receiving and transmitting of VLAN packets on a NIC even if no VLAN is registered with the card.

BZ#988460

When a slave device started up, the `current_arp_slave` parameter was unset but the active flags on the slave were not marked inactive. Consequently, more than one slave device with active flags in active-backup mode could be present on the system. A patch has been applied to fix this problem by marking the active flags inactive for a slave device before the `current_arp_slave` parameter is unset.

BZ#883575

Due to a bug in descriptor handling, the ioat driver did not correctly process pending descriptors on systems with the Intel Xeon Processor E5 family. Consequently, the CPU was utilized excessively on these systems. A patch has been applied to the ioat driver so the driver now determines pending descriptors correctly and CPU usage is normal again for the described processor family.

BZ#905561

A previous change in the bridge multicast code allowed sending general multicast queries in order to achieve faster convergence on startup. To prevent interference with multicast routers, send packets contained a zero source IP address. However, these packets interfered with certain multicast-aware switches, which resulted in the system being flooded with the IGMP membership queries with zero source IP address. A series of patches addresses this problem by disabling multicast queries by default and implementing multicast querier that allows to toggle up sending of general multicast queries if needed.

BZ#882413

A bug was causing bad block detection to try to isolate which blocks were bad in a device that had suffered a complete failure - even when bad block tracking was not turned on. This was causing very large delays in returning I/O errors when the entire set of RAID devices was lost to failure. The large delays caused problems during disaster recovery scenarios. The bad block tracking code is now properly disabled and errors return in a timely fashion when enough devices fail in a RAID array to exceed its redundancy.

BZ#876600

Previously, running commands such as "ls", "find" or "move" on a MultiVersion File System (MVFS) could cause a kernel panic. This happened because the `d_validate()` function, which is used for dentry validation, called the `kmem_ptr_validate()` function to validate a pointer to a parent dentry. The pointer could have been freed anytime so the `kmem_ptr_validate()` function could not guarantee the pointer to be dereferenced, which could lead to a NULL pointer dereference. This update modifies `d_validate()` to verify the parent-child relationship by traversing the parent dentry's list of child dentries, which solves this problem. The kernel no longer panics in the described scenario.

BZ#1008705

The sfc driver exposes on-board flash partitions using the MTD subsystem and it must expose up to 9

flash partitions per board. However, the MTD subsystem in Red Hat Enterprise Linux 6 has a static limit of 32 flash partitions. As a consequence, the Solarflare tools cannot operate on all boards if more than 3 boards are installed, preventing firmware on some boards from being updated or queried for a version number. With this update, a new `EFX_MCDI_REQUEST` sub-command has been added to the driver-private `SIOCEFX ioctl`, which allows bypassing the MTD layer and sending requests directly to the controller's firmware. The Solarflare tools can now be used and the firmware on all installed devices can be updated as expected in this scenario.

BZ#871795

Previously, the VLAN code incorrectly cleared the timestamping interrupt bit for network devices using the `igb` driver. Consequently, timestamping failed on the `igb` network devices with Precision Time Protocol (PTP) support. This update modifies the `igb` driver to preserve the interrupt bit if interrupts are disabled.

BZ#869736

When using more than 4 GB of RAM with an AMD processor, reserved regions and memory holes (E820 regions) can also be placed above the 4 GB range. For example, on configurations with more than 1 TB of RAM, AMD processors reserve the 1012 GB - 1024 GB range for the Hyper Transport (HT) feature. However, the Linux kernel does not correctly handle E820 regions that are located above the 4 GB range. Therefore, when installing Red Hat Enterprise Linux on a machine with an AMD processor and 1 TB of RAM, a kernel panic occurred and the installation failed. This update modifies the kernel to exclude E820 regions located above the 4 GB range from direct mapping. The kernel also no longer maps the whole memory on boot but only finds memory ranges that are necessary to be mapped. The system can now be successfully installed on the above-described configuration.

BZ#867689

The kernel interface to ACPI had implemented error messaging incorrectly. The following error message was displayed when the system had a valid ACPI Error Record Serialization Table (ERST) and the `pstore.backend` kernel parameter had been used to disable use of ERST by the `pstore` interface:

```
ERST: Could not register with persistent store
```

However, the same message was also used to indicate errors precluding registration. A series of patches modifies the relevant ACPI code so that ACPI now properly distinguish between different cases and accordingly prints unique and informative messages.

BZ#965132

When setting up a bonding device, a certain flag was used to distinguish between TLB and ALB modes. However, usage of this flag in ALB mode allowed enslaving NICs before the bond was activated. This resulted in enslaved NICs not having unique MAC addresses as required, and consequent loss of "reply" packets sent to the slaves. This patch modifies the function responsible for the setup of the slave's MAC address so the flag is no longer needed to discriminate ALB mode from TLB and the flag was removed. The described problem no longer occur in this situation.

BZ#920752

A bug in the `do_filp_open()` function caused it to exit early if any write access was requested on a read-only file system. This prevented the opening of device nodes on a read-only file system. With this update, the `do_filp_open()` has been fixed to no longer exit if a write request is made on a read-only file system.

BZ#981741

A dentry leak occurred in the FUSE code when, after a negative lookup, a negative dentry was neither dropped nor was the reference counter of the dentry decremented. This triggered a `BUG()` macro when unmounting a FUSE subtree containing the dentry, resulting in a kernel panic. A series of patches related to this problem has been applied to the FUSE code and negative dentries are now properly dropped so that triggering the `BUG()` macro is now avoided.

BZ#924804

This update reverts two previously-included `qla2xxx` patches. These patches changed the fibre channel target port discovery procedure, which resulted in some ports not being discovered in some corner cases. Reverting these two patches fixes the discovery issues.

BZ#957821

Due a bug in the memory mapping code, the `fdadvise64()` system call sometimes did not flush all the relevant pages of the given file from cache memory. A patch addresses this problem by adding a test condition that verifies whether all the requested pages were flushed and retries with an attempt to empty the LRU pagevecs in the case of test failure.

BZ#957231

The `xen-netback` and `xen-netfront` drivers cannot handle packets with size greater than 64 KB including headers. The `xen-netfront` driver previously did not account for any headers when determining the maximum size of GSO (Generic Segmentation Offload). Consequently, Xen DomU guest operations could have caused a network DoS issue on DomU when sending packets larger than 64 KB. This update adds a patch that corrects calculation of the GSO maximum size and the problem no longer occurs.

BZ#848085

A possible race in the `tty` layer could result in a kernel panic after triggering the `BUG_ON()` macro. As a workaround, the `BUG_ON()` macro has been replaced by the `WARN_ON()` macro, which allows for avoiding the kernel panic and investigating the race problem further.

BZ#980876

A bug in the network bridge code allowed an internal function to call code which was not atomic-safe while holding a spin lock. Consequently, a "BUG: scheduling while atomic" error has been triggered and a call trace logged by the kernel. This update applies a patch that orders the function properly so the function no longer holds a spin lock while calling code which is not atomic-safe. The aforementioned error with a call trace no longer occurs in this case.

BZ#916806

An NFSv4 client could previously enter a deadlock situation with the state recovery thread during state recovery after a reboot of an NFSv4 server. This happened because the client did not release the NFSv4 sequence ID of an `OPEN` operation that was requested before the reboot. This problem is resolved by releasing the sequence ID before the client starts waiting for the server to recover.

BZ#859562

A bug in the device-mapper RAID kernel module was preventing the "sync" directive from being honored. The result was that users were unable to force their RAID arrays to undergo a complete resync if desired. This has been fixed and users can use `lvchange --resync my_vg/my_raid_lv` to force a complete resynchronization on their LVM RAID arrays.

Enhancements

BZ#823012

This update provides simplified performance analysis for software on Linux on System z by using the Linux perf tool to access the hardware performance counters.

BZ#829506

The fnic driver previously allowed I/O requests with the number of SGL descriptors greater than is supported by Cisco UCS Palo adapters. Consequently, the adapter returned any I/O request with more than 256 SGL descriptors with an error indicating invalid SGLs. A patch has been applied to limit the maximum number of supported SGLs in the fnic driver to 256 and the problem no longer occurs.

BZ#840454

To transmit data, for example, trace data, from guests to hosts, a low-overhead communication channel was required. Support for the splice() call has been added to the virtio_console module in the Linux kernel. This enables sending guest kernel data to the host without extra copies of the data being made inside the guest. Low-overhead communication between the guest Linux kernel and host userspace is performed via virtio-serial.

BZ#888903

A new MTIOCTOP operation, MTWEOF1, has been added to the SCSI tape driver, which allows writing of "filemarks" with the "immediate" bit. This allows a SCSI tape drive to preserve the content of its buffer, enabling the next file operation to start immediately. This can significantly increase write performance for applications that have to write multiple small files to the tape while it also reduces tape weariness.

BZ#913650

Previously, a user needed to unmount, deactivate their RAID LV, and re-activate it in order to restore a transiently failed device in their array. Now it is possible to restore such devices without unmounting by simply running 'lvchange --refresh'.

BZ#923212

Open vSwitch (OVS) is an open-source, multi-layer software switch designed to be used as a virtual switch in virtualized server environments. Starting with Red Hat Enterprise Linux 6.4, the Open vSwitch kernel module is included as an enabler for Red Hat Enterprise Linux OpenStack Platform. Open vSwitch is only supported in conjunction with Red Hat products containing the accompanying user-space packages. Without these packages, Open vSwitch will not function and cannot be used with other Red Hat Enterprise Linux variants.

BZ#928983

The RHEL6.5 bfa driver changes behavior of the dev_loss_tmo value such that it can only be set to a value greater than the bfa driver specific path_tov value. The minimum default value that the dev_loss_tmo can be set to is 31 seconds. Attempting to set the dev_loss_tmo value lower than 31 seconds without lowering the default bfa path_tov value will not succeed.

BZ#929257

Error recovery support has been added to the flash device driver, which allows hardware service upgrades without negative impact on I/O of flash devices.

BZ#929259

The crypto adapter resiliency feature has been added. This feature provides System z typical RAS for cryptographic adapters through comprehensive failure recovery. For example, this feature handles unexpected failures or changes caused by Linux guest relocation, suspend and resume activities or configuration changes.

BZ#929262

The "fuzzy live dump" feature has been added. With this feature kernel dumps from running Linux systems can be created, to allow problem analysis without taking down systems. Because the Linux system continues running while the dump is written, and kernel data structures are changing during the dump process, the resulting dump contains inconsistencies.

BZ#929264, BZ#929264

The kernel now provides an offline interface for DASD devices. Instead of setting a DASD device offline and returning all outstanding I/O requests as failed, with this interface you can set a DASD device offline and write all outstanding data to the device before setting the device offline.

BZ#929274

The kernel now provides the Physical Channel ID (PCHID) mapping that enables hardware detection with a machine-wide unique identifier.

BZ#929275

The kernel now provides VEPA mode support. VEPA mode routes traffic between virtual machines on the same mainframe through an external switch. The switch then becomes a single point of control for security, filtering, and management.

BZ#755486, BZ#755486

Message Transfer Part Level 3 User Adaptation Layer (M3UA) is a protocol defined by the IETF standard for transporting MTP Level 3 user part signaling messages over IP using Stream Control Transmission Protocol (SCTP) instead of telephony equipment like ISDN and PSTN. With this update, M3UA measurement counters have been included for SCTP.

BZ#818344

Support for future Intel 2D and 3D graphics has been added to allow systems using future Intel processors to be certified through the Red Hat Hardware Certification program.

BZ#826061

In certain storage configurations (for example, configurations with many LUNs), the SCSI error handling code can spend a large amount of time issuing commands such as TEST UNIT READY to unresponsive storage devices. A new sysfs parameter, `eh_timeout`, has been added to the SCSI device object, which allows configuration of the timeout value for TEST UNIT READY and REQUEST SENSE commands used by the SCSI error handling code. This decreases the amount of time spent checking these unresponsive devices. The default value of `eh_timeout` is 10 seconds, which was the timeout value used prior to adding this functionality.

BZ#839470, BZ#839470

With this update, 12Gbps LSI SAS devices are now supported in Red Hat Enterprise Linux 6.

BZ#859446

Red Hat Enterprise Linux 6.5 introduces the Orlov block allocator that provides better locality for files which are truly related to each other and likely to be accessed together. In addition, when resource groups are highly contended, a different group is used to maximize performance.

BZ#869622

The mdadm tool now supports the TRIM commands for RAID0, RAID1, RAID10 and RAID5.

BZ#880142

Network namespace support for OpenStack has been added. Network namespaces (netns) is a lightweight container-based virtualization technology. A virtual network stack can be associated with a process group. Each namespace has its own loopback device and process space. Virtual or real devices can be added to each network namespace, and the user can assign IP addresses to these devices and use them as a network node.

BZ#908606

Support for dynamic hardware partitioning and system board slot recognition has been added. The dynamic hardware partitioning and system board slot recognition features alert high-level system middleware or applications for reconfiguration and allow users to grow the system to support additional workloads without reboot.

BZ#914771, BZ#920155, BZ#914797, BZ#914829, BZ#914832, BZ#914835

An implementation of the Precision Time Protocol (PTP) according to IEEE standard 1588 for Linux was introduced as a Technology Preview in Red Hat Enterprise Linux 6.4. The PTP infrastructure, both kernel and user space, is now fully supported in Red Hat Enterprise Linux 6.5. Network driver time stamping support now also includes the following drivers: bnx2x, tg3, e1000e, igb, ixgbe, and sfc.

BZ#862340

The Solarflare driver (sfc) has been updated to add PTP support as a Technology Preview.

BZ#918316

In Red Hat Enterprise Linux 6.5, users can change the cryptography hash function from MD5 to SHA1 for Stream Control Transmission Protocol (SCTP) connections.

BZ#922129

The pm8001/pm80xx driver adds support for PMC-Sierra Adaptec Series 6H and 7H SAS/SATA HBA cards as well as PMC Sierra 8081, 8088, and 8089 chip based SAS/SATA controllers.

BZ#922299

VMware Platform Drivers Updates The VMware network para-virtualized driver has been updated to the latest upstream version.

BZ#922941

The Error-correcting code (ECC) memory has been enabled for future generation of AMD processors. This feature provides the ability to check for performance and errors by accessing ECC memory related counters and status bits.

BZ#922965

Device support is enabled in the operating system for future Intel System-on-Chip (SOC) processors. These include Dual Atom processors, memory controller, SATA, Universal Asynchronous

Receiver/Transmitter, System Management Bus (SMBUS), USB and Intel Legacy Block (ILB - lpc, timers, SMBUS (i2c_801 module)).

BZ#947944

Kernel Shared Memory (KSM) has been enhanced to consider non-uniform memory access (NUMA) when coalescing pages, which improves performance of the applications on the system. Also, additional page types have been included to increase the density of applications available for Red Hat OpenShift.

BZ#949805

FUSE (Filesystem in User Space) is a framework that allows for development of file systems purely in the user space without requiring modifications to the kernel. Red Hat Enterprise Linux 6.5 delivers performance enhancements for user space file systems that use FUSE, for example, GlusterFS (Red Hat Storage).

BZ#864597

The default TCP stack buffers are too large for high bandwidth applications that fully utilize the Ethernet link. This could result in a situation where connection bandwidth could not be fully utilized and could be distributed unequally if the link was shared by multiple client devices. To resolve this problem, a new feature, TCP Small Queues (TSQ), has been introduced to the TCP code. The TSQ feature reduces a number of TCP packets in xmit queues, TCP round-trip time (RTT), and the congestion window (CWND) size. It also mitigates an impact of a possible bufferbloat problem. This change also includes a patch that resolves a performance problem on mlx4 devices caused by setting the default value of the Tx coalescing too high.

All Red Hat Enterprise Linux 6 users are advised to install these updated packages, which correct these issues, and fix the bugs and add the enhancements noted in the Red Hat Enterprise Linux 6.5 Release Notes and Technical Notes. The system must be rebooted for this update to take effect.

8.83. KEXEC-TOOLS

8.83.1. RHBA-2013:1576 — kexec-tools bug fix and enhancement update

Updated kexec-tools packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The kexec-tools packages contain the /sbin/kexec binary and utilities that together form the user-space component of the kernel's kexec feature. The /sbin/kexec binary facilitates a new kernel to boot using the kernel's kexec feature either on a normal or a panic reboot. The kexec fastboot mechanism allows booting a Linux kernel from the context of an already running kernel.

Bug Fixes

BZ#1015764

Previously, in the mkdumprd utility, the strip_comments() function was not implemented correctly. When arguments were passed to strip_comments(), it only took the first argument into account and skipped the rest. As a consequence, it passed the "makedumpfile" argument to the \$config_val variable, but the parameters for "makedumpfile" were missed. With this update, the strip_comments() function has been modified. As a result, it no longer skips arguments passed to it.

BZ#886572

When the kdump file system resided on a logical volume or a volume group with another independent and encrypted device, the mkdumprd utility exited with an error message when trying to access the encrypted device, preventing kdump from functioning properly. A patch has been provided to address this problem and kdump is now properly reconfigured and restarted in the described scenario, thus fixing this bug.

BZ#920705

Certain multi-port network cards return the same PCI bus address for all ports. When the kdump utility maps the network ports, it cannot differentiate one network port from another on these cards. Consequently, when different network ports were on different networks, kdump failed to dump data over NFS or SSH. This update ensures that the MAP_NET_BY_MAC variable is set in the described scenario and kdump now dumps data for all ports as expected.

BZ#883543

Previously, a udev rule in the 98-kexec.rules file spawned processes that restarted the kdump tool with each memory added. To fix this bug, the "condrestart" parameter is used when attempting to restart a service that was previously running. As a result, kdump is no longer restarted when a restart is not needed.

BZ#921142

Previously, kernel modules in the extra_modules list were overridden by the built-in blacklist. Consequently, kdump was unable to load the mlx4_core and mlx4_en modules and dump data over network cards using these modules. With this update, modules in the extra_modules list are not excluded if they are blacklisted and kdump can use them as expected.

BZ#1008543

Previously, in makedumpfile, the dumpfile header had a field which was inherited from the deprecated "diskdump" facility. The field was used by the crash utility as a delimiter to determine whether a physical address read request was legitimate. The field could not handle Physical Frame Number (PFN) values greater than 32-bits and such values were truncated. This update adds three new fields to the header. As a result, the dumpfile header in makedumpfile correctly handles PFN values greater than 32-bits.

BZ#876667

Previously, for some kernel modules, the "modprobe --show-depends" command's output did not have the "insmod" prefix for every line. Consequently, the mkdumprd utility failed to load as the current code assumed that each line started with the "insmod" prefix. The code has been modified to only match lines starting with "insmod" in awk scripts. As a result, mkdumprd no longer fails to load in this scenario.

BZ#1009207

Previously, in cyclic mode, the makedumpfile recalculated incorrectly the size of the cyclic buffer size. As a consequence, makedumpfile did not update the length of the range of a cycle in page frame numbers, which caused a buffer overrun or a segmentation violation. Furthermore, due to the divideup() function in the recalculations, the cyclic buffer size became too much aligned and less efficient. A patch has been provided to fix these bugs and the aforementioned problems no longer occur in this scenario.

BZ#1010103

The x86_64 kernel is a relocatable kernel, and there can be a gap between the physical address statically assigned to the kernel data and texts, and the address that is really assigned to each object corresponding to the kernel symbols. The gap is the phys_base() function. The makedump utility

calculates the `phys_base` in an ad-hoc way that compares the addresses of some of occurrences of "Linux kernel" strings in certain range of the `vmcore`. As a consequence, `makedumpfile` failed calculating `phys_base` and also failed converting a `vmcore`. This bug has now been fixed and `makedumpfile` calculates `phys_base` correctly and converts `vmcore` normally.

BZ#893764

Previously, setting empty Direct Access Storage Device (DASD) options, parsed from the `/etc/dasd.conf` file, resulted in displaying environment variables. As a consequence, restarting the `kdump` service displayed the complete `kdump` script. After this update, if there are no options specified in the `/etc/dasd.conf` file for a device, the `kdump` script proceeds to the next one. As a result, restarting the `kdump` service no longer displays the complete `kdump` script.

BZ#918372

Previously, `kdump` data written on a raw device was not completely flushed. As a consequence, the saved `vmcore` was occasionally incomplete. This update uses the `blockdev` tool to flush out block device buffers. As a result, `vmcore` saved on a raw device is now always complete.

BZ#903529

Previously, because Storage Class Memory (SCM) devices did not expose the same `sysfs` attributes as Small Computer System Interface (SCSI) disks, the `mkdumprd` utility failed to determine the list of "critical disks" for writing a dump file. As a consequence, certain SCM devices were not correctly handled by `mkdumprd`, resulting in an infinite loop when trying to specify a file system on such a device as target for `kdump`. After this update, `mkdumprd` now handles waiting for SCM devices based on the device's storage increment address, a property which uniquely identifies an SCM device across reboots. As a result, `mkdumprd` now successfully determines the list of "critical disks" for writing a dump file and an infinite loop no longer occurs.

BZ#906601

Previously, on a system configured with multipath support, the `mkdumprd` tool pushed the code handling multipath devices into the `kdump` `initrd`. As a consequence, the `kdump` utility failed to capture `vmcore` on multipath devices. This update introduces a mechanism where the call to the `kpartx` utility is delayed until the "`dmsetup ls`" command lists the device names which match the multipath device where a `vmcore` is going to be captured. As a result, `mkdumprd` now waits until the multipath devices are created and then successfully captures a `vmcore` on them.

BZ#977651

Previously, when Red Hat Enterprise Linux was configured to use the `hugepages` parameter, the `kdump` kernel also used this parameter. As a consequence, due to its limited memory, using `hugepages` could lead to an Out Of Memory (OOM) error for the `kdump` kernel. With this update, `hugepages` and `hugepagesz` kernel parameters are not used by the `kdump` kernel when the Red Hat Enterprise Linux is using them. If the user wants to explicitly use `hugepages` in the `kdump` kernel, they can be specified through the `KERNEL_COMMANDLINE_APPEND` option in the `/etc/sysconfig/kdump` file.

BZ#963948

Previously, when a VMware guest was added additional RAM, multiple instances of the `kdump.init` script were started concurrently. As a consequence, a race condition occurred among the `kdump.init` instances. By introducing a global mutex lock, now only one instance can acquire this lock and run, others will be waiting for the lock in queue. As a result, the `kdump.init` instances are run in serial order and a race condition no longer occurs in this scenario.

BZ#951035

Previously, when the `e2fsprogs` package, which contains tools used by the `mkdumprd` utility, was not installed on the system, `mkdumprd` displayed a misleading error message. With this update, the error message has been improved to explicitly inform the user which of these tools is missing.

Enhancements

BZ#959449

This update allows the `kdump` tool to work with an arbitrary bridge, bond or vlan names over a network. Now, it is possible to name a device without following the established naming conventions, for example, bonding devices do not need to start with "bond". The user can determine whether a network device is a bond, bridge or vlan by checking for the existence of specific directories in the `/sys/` or `/proc/` directories.

BZ#871522

With this update, `kexec-tools` now respect the memory limit while building crash memory ranges on a 64-bit PowerPC. The kernel exports memory limit information through the `/proc/device-tree` file, which `kexec-tools` now read and limit the crash memory ranges accordingly.

BZ#825476, BZ#902147, BZ#902148

In Red Hat Enterprise Linux 6.5, the `makedumpfile` utility supports the Lempel–Ziv–Oberhumer (LZO) and snappy compression formats. Using these compression formats instead of the `zlib` format is quicker, in particular when compressing data with randomized content.

BZ#947621

This update includes changes to allow filtering of poisoned pages during a crash dump capture. The user can now decide if poisoned pages are dumped. Furthermore, filtering can increase dumping speed.

BZ#797231

This update adds an SELinux relabeling during `kdump` service startup. The `kdump` service now relabels files in the dumping path which have an incorrect or missing label.

BZ#909402

In previous Red Hat Enterprise Linux releases, support for SSH FIPS mode was incomplete. This update adds the relevant library files and `*.hmac` files to the `kdump` kernel. The `kdump` utility can now work in SSH FIPS mode.

BZ#975642

This update adds documentation for the `--allow-missing` `mkdumprd` option to the `mkdumprd(8)` manual page.

Users of `kexec-tools` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.84. KSH

8.84.1. **RHBA-2013:1599 — ksh bug fix and enhancement update**

Updated ksh packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

KornShell (KSH) is a Unix shell developed by AT&T Bell Laboratories, which is backward-compatible with the Bourne shell (Bash) and includes many features of the C shell. The most recent version is KSH-93. KornShell complies with the POSIX.2 standard (IEEE Std 1003.2-1992).



NOTE

The ksh package has been upgraded to upstream version 20120801, which provides a number of bug fixes and enhancements over the previous version. (BZ#[840568](#))

Bug Fixes

BZ#[761551](#)

Previously, the ksh shell did not set any editing mode as default, which caused various usability problems in interactive mode and with shell auto-completion. This update sets emacs editing mode as default for new users. As a result, the usability is significantly improved and the shell auto-completion works as expected.

BZ#[858263](#)

Previously, the ksh internal counter of jobs was too small. Consequently, when a script used a number of subshells in a loop, a counter overflow could occur causing the ksh shell to terminate unexpectedly with a segmentation fault. This update modifies ksh to use bigger types for counter variables. As a result, ksh no longer crashes in the described scenario.

BZ#[903750](#)

Previously, the ksh shell did not compute an offset for fixed size variables correctly. As a consequence, when assigning a right-justified variable with a fixed width to a smaller variable, the new variable could have an incorrect content. This update applies a patch to fix this bug and the assignment now proceeds as expected.

BZ#[913110](#)

Previously, the output of command substitutions was not always redirected properly. Consequently, the output in a here-document could be lost. This update fixes the redirection code for command substitutions and the here-document now contains the output as expected.

BZ#[921455](#), BZ#[982142](#)

Using arrays inside of ksh functions, command aliases, or automatically loaded functions caused memory leaks to occur. The underlying source code has been modified to fix this bug and the memory leaks no longer occur in the described scenario.

BZ#[922851](#)

Previously, the ksh SIGTSTP signal handler could trigger another SIGTSTP signal. Consequently, ksh could enter an infinite loop. This updated version fixes the SIGTSTP signal processing and ksh now handles the signal without any problems.

BZ#[924440](#)

Previously, the ksh shell did not resize the file descriptor list every time it was necessary. This could lead to memory corruption when several file descriptors were used. As a consequence, ksh terminated unexpectedly. This updated version resizes the file descriptor list every time it is needed,

and ksh no longer crashes in the described scenario.

BZ#960034

Previously, the ksh shell ignored the "-m" argument specified by the command line. As a consequence, ksh did not enable monitor mode and the user had to enable it in a script. With this update, ksh no longer ignores the argument so that the user is able to enable monitor mode from the command line as expected.

BZ#994251

The ksh shell did not handle I/O redirections from command substitutions inside a pipeline correctly. Consequently, the output of certain commands could be lost. With this update, the redirections have been fixed and data is no longer missing from the command outputs.

Users of ksh are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.85. LEDMON

8.85.1. RHBA-2013:1722 — ledmon bug fix and enhancement update

Updated ledmon packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The ledmon and ledctl utilities are user-space applications designed to control LEDs associated with each slot in an enclosure or a drive bay. There are two types of systems: 2-LED system (Activity LED, Status LED) and 3-LED system (Activity LED, Locate LED, Fail LED). Users must have root privileges to use this application.

**NOTE**

The ledmon packages have been upgraded to upstream version 0.78, which provides a number of bug fixes and enhancements over the previous version. (BZ#922976, BZ#876593, BZ#887370)

Users of ledmon are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.86. LIBXCURSOR

8.86.1. RHBA-2013:0906 — libXcursor bug fix update

Updated libXcursor packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The X.Org's X11 libXcursor package provides a runtime library for cursor management, designed to help locate and load cursors.

Bug Fix**BZ#949586**

In the last rebuild of libXcursor, the Icon Theme was changed to Adwaita, which was not available in Red Hat Enterprise Linux 6. To fix this bug, the Icon Theme has been changed back to dmz-aa for Red Hat Enterprise Linux 6.

Users of libXcursor are advised to upgrade to these updated packages, which fix this bug.

8.87. LIBCGROUP

8.87.1. RHBA-2013:1685 — libcgroup bug fix and enhancement update

Updated libcgroup packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The libcgroup packages provide tools and libraries to manage and monitor control groups.

Bug Fixes

BZ#972893

Previously, the pam_cgoup pluggable authentication module (PAM) did not use caching. As a consequence, when a system had several thousand users and the cgrules.conf file contained several thousand lines of configuration settings, the login time could take several seconds. With this update, the libcgroup code no longer reads the /etc/passwd file once for every line in cgrules.conf, and the login time is no longer affected in the described scenario.

BZ#863172

Prior to this update, the cgroup files did not have write permissions set correctly. Consequently, members of the group that owned the cgroup files could not modify their content. The group permissions have been updated, and the members of the group can now modify the content of the cgroup files.

BZ#921328

Previously, the behavior of the cgred service when opening the configuration file was not set correctly. Consequently, cgred failed to start if the configuration file was missing or empty. Explicit checks for the existence of the configuration file have been removed, and cgred now starts with a missing or empty configuration file as expected.

BZ#912425

The code in the cg_get_pid_from_flags() function assumed that every entry in the /etc/cgrules.conf file had the process name specified. As a consequence, if the entry in the /etc/cgrules.conf file did not specify the process name, the cgred service terminated unexpectedly with a segmentation fault. This update allows the code to accept empty process names and cgred no longer crashes.

BZ#946953

Prior to this update, the permissions of the /bin/cgclassify file were set incorrectly. As a consequence, the "--sticky" option of the cgclassify command was ignored when running under a non-privileged user. The file permissions of /bin/cgclassify have been updated, and the "--sticky" option now works correctly for regular users.

BZ#753334

Previously, using commas in the lexical analyzer was not supported. As a consequence, the cgconfig service failed to parse commas in the cgconfig.conf file. Support for commas in the lexical analyzer has been added, and cgconfig can now successfully parse commas in cgconfig.conf.

BZ#924399

The cgrulesengd daemon had different default logging level than the rest of the library. Consequently, the log messages were inconsistent. With this update, the logging level of the cgrulesengd daemon and the library has been unified, and the log messages are now consistent as expected.

BZ#809550

Prior to this update, the cgcreate(1) manual page contained the invalid "-s" option in the synopsis. This update removes this option.

BZ#961844

Previously, the cgraded service was starting too early in the boot process. As a consequence, if some services started before cgraded, they could avoid being restricted. The boot priority of cgraded has been lowered, and all services are now restricted correctly.

Enhancement**BZ#589535**

After this update, the cgraded daemon supports automated control groups for every user in any UNIX group that logs in. A template is now used to create a new control group automatically, and every process the user launches is started in the appropriate group, which makes managing multiple users easier.

Users of libcgroup are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.88. LIBDRM

8.88.1. [RHBA-2013:1551 — libdrm bug fix and enhancement update](#)

Updated libdrm packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Direct Rendering Manager runtime library (libdrm) provides a user-space interface library for direct rendering clients.

**NOTE**

The libdrm packages have been upgraded to upstream version 2.4.45, which provides a number of bug fixes and enhancements over the previous version. (BZ#914774)

Users of libdrm are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.89. LIBGUESTFS

8.89.1. RHSA-2013:1536 — Moderate: libguestfs security, bug fix, and enhancement update

Updated libguestfs packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE link(s) associated with each description below.

Libguestfs is a library and set of tools for accessing and modifying guest disk images.



NOTE

The libguestfs package has been upgraded to upstream version 1.20.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#958183)

Security Fix

CVE-2013-4419

It was found that `guestfish`, which enables shell scripting and command line access to libguestfs, insecurely created the temporary directory used to store the network socket when started in server mode. A local attacker could use this flaw to intercept and modify other user's `guestfish` command, allowing them to perform arbitrary `guestfish` actions with the privileges of a different user, or use this flaw to obtain authentication credentials.

This issue was discovered by Michael Scherer of the Red Hat Regional IT team.

Bug Fixes

BZ#892291

Previously, when the `guestmount` utility failed to create a hard link, an incorrect error message was returned. Consequently, information about the true cause of error was not displayed. With this update, the error handling in `guestmount` has been fixed and correct messages are now displayed in the described case.

BZ#892834

When attempting to rename a symbolic link with the `guestmount` utility, `guestmount` followed the link instead of overwriting it. With this update, a `guestfs_rename` API has been added, which allows `guestmount` to rename target files correctly.

BZ#908255

Downloading a directory using the `guestfs_download` API or the `guestfish download` command is not allowed. However `libguestfs` did not return an error in such case and lost protocol synchronization instead. With this update, `libguestfs` now tests if the download source is a directory and returns an error message if it is.

BZ#909666

Under certain circumstances, long-running libguestfs API calls, which generated progress messages, caused libguestfs to terminate unexpectedly due to a stack overflow. The underlying source code has been modified to handle this case and the stack overflow no longer occurs.

BZ#971090

Prior to this update, the `libguestfs` inspection did not detect a Microsoft Windows guest that used a non-standard `systemroot` path. With this update, `libguestfs` has been modified to use the contents of the Windows `boot.ini` file to find the `systemroot` path. As a result, Windows guests are detected properly even if they use non-standard `systemroot` paths.

BZ#971326

Previously, `libguestfs` did not resize a Microsoft Windows NTFS file system when the target size was not explicitly specified. With this update, `libguestfs` has been modified to establish this size automatically from the target storage device. As a result, NTFS file systems can now be resized even without specifying the target size.

BZ#975753

The `virt-resize` fails on Windows guests that are in an inconsistent state. This update adds the description of this problem to the `guestfs(3)` man page.

BZ#975760

If the `iface` parameter was used when adding a drive, `libguestfs` entered an infinite loop. With this update, `libguestfs` has been fixed to process `iface` parameters correctly, thus preventing the hang.

BZ#980358

Calling the `guestfs_filesystem_available(g, "xfs")` function could be evaluated as true even if certain XFS functions were not available. This problem has been documented in the `guestfs(3)` man page.

BZ#980372

Prior to this update, the `hivex-commit` command with a relative path parameter wrote to a location inaccessible to users. This command has been modified to require an absolute path or a NULL path that overwrites the original. An error message is now displayed if a relative path is passed to `hivex-commit`.

BZ#985269

The syntax for setting Access Control Lists (ACLs) with `libguestfs` is now documented in the `guestfs(3)` man page.

BZ#989352

When `libguestfs` was used to read the capabilities of a file that had no capabilities set, `libguestfs` returned an error. The `guestfs_cap_get_file()` function that is responsible for retrieving the file capabilities has been modified to return an empty string in the described case.

BZ#996039

Under certain circumstances, using the `guestfish` command with both `--remote` and `--add` options can have unexpected results. This behavior has been documented in the `guestfish(1)` man page.

BZ#996825

Previously, when using the `guestfish --remote` command, the following message was displayed:

```
libguestfs: error: waitpid (qemu): No child processes
```

With this update, this unnecessary message is no longer displayed.

BZ#998108

Previously, when the `libguestfs` package was used on systems under heavy load, messages about "unstable clocks" appeared in the debugging output. With this update, `libguestfs` has been modified to check if the `kvmclock` kernel feature is enabled, thus reducing the aforementioned message output.

BZ#1000122

Prior to this update, using the `guestfs_sh` or `sh` command before mounting a disk caused the `guestfish` utility to terminate with a segmentation fault. With this update, `guestfish` has been modified to verify if a file system is mounted before executing these commands, and if not, an error message is displayed. As a result, `guestfish` no longer crashes in the aforementioned scenario.

Users of `libguestfs` are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.90. LIBIBVERBS-ROCEE

8.90.1. RHEA-2013:1740 — libibverbs-rocee and libmlx4-rocee bug fix and enhancement update

Updated `libibverbs-rocee` and `libmlx4-rocee` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Red Hat Enterprise Linux includes a collection of InfiniBand and iWARP utilities, libraries, and development packages for writing applications that use Remote Direct Memory Access (RDMA) technology.



NOTE

The `libibverbs-rocee` packages have been upgraded to upstream version 1.1.7 and the `libxml-rocee` packages to upstream version 1.0.5, which provides a number of bug fixes and enhancements over the previous versions and keeps the HPN channel synchronized with the base Red Hat Enterprise Linux channel, where the sister versions of these packages (`libibverbs` and `libmlx4`) were also updated to the latest upstream release.

All users of Remote Direct Memory Access (RDMA) technology are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.91. LIBKSBA

8.91.1. RHBA-2013:0837 — libksba bug fix update

Updated `libksba` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

KSBA is a library designed to build software based on the X.509 and CMS standards. It provides developers with a single API that handles the underlying details of the X.509 standard and presents data consistently.

Bug Fix

BZ#658058

Previously, contents of the `/usr/bin/libksba-config` script conflicted between 32-bit and 64-bit versions of `libksba-devel` packages. Consequently, these packages could not be installed simultaneously. This update amends the script to make its contents consistent for all architectures, thus fixing this bug.

Users of `libksba` are advised to upgrade to these updated packages, which fix this bug.

8.92. LIBNL

8.92.1. RHBA-2013:1730 — libnl bug fix update

Updated `libnl` packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The `libnl` packages contain a convenience library to simplify using the Linux kernel's Netlink sockets interface for network manipulation.

Bug Fixes

BZ#682240

When a domain was started using the `libvirt` client libraries and utilities, a memory leak was triggered from the `libnl` library because `libnl` continued to use memory that was no longer in use. With this update, memory leaks in `libnl` are fixed, and `libnl` releases memory after it completes its usage.

BZ#689559

Prior to this update, `libnl`'s error handling made generous use of the `strerror()` function. Nevertheless, the `strerror()` function was not threadsafe, and it was possible for multiple threads in an application to call `libnl`. With this update, all the occurrences of `strerror()` are replaced with a call to the `strerror_r()` function that puts the message into a thread-local static buffer.

BZ#953339

When the `max_vfs` parameter of the `igb` module, which allocates the maximum number of Virtual Functions, was set to any value greater than 50,50 on a KVM (Kernel-based Virtual Machine) host, the guest failed to start with the following error messages:

```
error : virNetDevParseVfConfig:1484 : internal error missing IFLA_VF_INFO in netlink response
```

```
error : virFileReadAll:457 : Failed to open file '/var/run/libvirt/qemu/eth0_vf0': No such file or directory
error : virFileReadAll:457 : Failed to open file '/var/run/libvirt/qemu/eth1_vf0': No such file or directory
```

This update increases the default receive buffer size to allow receiving of Netlink messages that exceed the size of a memory page. Thus, guests are able to start on the KVM host, and error messages no longer occur in the described scenario.


Users of `libnl` are advised to upgrade to these updated packages, which fix these bugs.

8.93. LIBPCAP

8.93.1. RHBA-2013:1727 — libpcap bug fix and enhancement update

Updated libpcap packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Packet Capture library (pcap) provides a high level interface to packet capture systems. All packets on the network, even those destined for other hosts, are accessible through this mechanism. It also supports saving captured packets to a 'savefile', and reading packets from a 'savefile'. libpcap provides implementation-independent access to the underlying packet capture facility provided by the operating system.



NOTE

The libpcap packages have been upgraded to upstream version 1.4.1, which provides a number of bug fixes and enhancements over the previous version. (BZ#[916749](#))

Bug Fixes

BZ#[723108](#)

Previously, the libpcap library generated wrong filtering code for Berkeley Packet Filter (BPF) infrastructure. As a consequence, the in-kernel packet filter was discarding some packets which should have been received by userspace process. Moreover, the tcpdump utility produced incorrect output when a fragmentation of IPv6 packet occurred because of the MTUlink. To fix this bug, the code which deals with BPF filter generation has been fixed to check for fragmentation headers in IPv6 PDUs before checking for the final protocol. As a result, the kernel filter no longer discards IPv6 fragments when source-site fragmentation occurs during IPv6 transmission and tcpdump receives all packets.

BZ#[731789](#)

Prior to this update, libpcap was unable to open a capture device with small values of SnapLen, which caused libpcap to return an error code and tcpdump to exit prematurely. Calculation of frames for memory mapping packet capture mechanism has been adjusted not to truncate packets to smaller values than actual SnapLen, thus fixing the bug. As a result, libpcap no longer returns errors when trying to open a capture device with small values of SnapLen, and applications using libpcap are able to process packets.

Users of libpcap are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.94. LIBQB

8.94.1. [RHBA-2013:1634](#) — libqb bug fix and enhancement update

Updated libqb packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The libqb packages provide a library with the primary purpose of providing high performance client server reusable features, such as high performance logging, tracing, inter-process communication, and polling.

**NOTE**

The libqb packages have been upgraded to upstream version 0.16.0, which provides a number of bug fixes and enhancements over the previous version, including a patch to fix a bug in the `qb_log_from_external_source()` function that caused the Pacemaker's policy engine to terminate unexpectedly. (BZ#950403)

Bug Fix**BZ#889299**

Output of the Blackbox window manager did not contain logging information if the string's length or precision was specified. This affected usability of the Blackbox output for debugging purposes, specifically when used with the Pacemaker cluster resource manager. The problem was caused by bugs in the libqb's implementation of the `strncpy()` and `strncat()` functions and the code responsible for the Blackbox log formatting. This update corrects these bugs so the Blackbox output is now formatted as expected.

Users of libqb are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.95. LIBREOFFICE

8.95.1. RHBA-2013:1594 — libreoffice bug fix and enhancement update

Updated libreoffice packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

LibreOffice is an Open Source, community-developed, office productivity suite. It includes the key desktop applications, such as a word processor, spreadsheet, presentation manager, formula editor and drawing program. LibreOffice replaces OpenOffice.org and provides a similar but enhanced and extended Office Suite.

**NOTE**

The libreoffice package has been upgraded to upstream version 4.0.4, which provides a number of bug fixes and enhancements over the previous version. (BZ#919230)

Bug Fixes**BZ#820554**

The `--enable-new-dtags` flag was added to allow certain types of built time regression tests to function. As a consequence, the GCJ Java compiler failed to search the correct location of Java libraries. This update applies a patch to remove the flag and GCJ works as expected.

BZ#829709

Previously, the LibreOffice suite was not fully translated into certain local languages. This update provides the full translation of LibreOffice to local languages.

BZ#833512

During upgrading the OpenOffice.org suite to the OpenOffice suite, backward compatibility links were removed and the OpenOffice.org icons were not migrated to LibreOffice. Consequently, an attempt to

launch LibreOffice failed with an error. With this update, the compatibility links have been restored and the icons now work as expected.

BZ#847519

Due to a bug in the chart creation code, an attempt to create a chart, under certain circumstances, failed with a segmentation fault. The underlying source code has been modified to fix this bug and the chart creation now works as expected.

BZ#855972

Due to a bug in the underlying source code, an attempt to show the outline view in the Impress utility terminated unexpectedly. This update applies a patch to fix this bug and the outline view no longer crashes in the described scenario.

BZ#863052

Certain versions of the Microsoft Office suite contain mismatching internal time stamp fields. Previously, the LibreOffice suite detected those fields and returned exceptions. Consequently, the user was not able to open certain Microsoft Office documents. With this update, LibreOffice has been modified to ignore the mismatching time stamp fields and Microsoft Office documents can be opened as expected.

BZ#865058

When a large amount of user-defined number formats was specified in a file, those formats used all available slots in a table and for remaining formats the general format was used. As a consequence, certain cell formatting did not preserve during loading the file. With this update, a patch has been provided and cell formatting works as expected.

BZ#871462

The Libreoffice suite contains a number of harmless files used for testing purposes. Previously, on Microsoft Windows system, these files could trigger false positive alerts on various anti-virus software, such as Microsoft Security Essentials. For example, the alerts could be triggered when scanning the Red Hat Enterprise Linux 6 ISO file. The underlying source code has been modified to fix this bug and the files no longer trigger false positive alerts in the described scenario.

BZ#876742

Due to an insufficient implementation of tables, the Impress utility made an internal copy of a table during every operation. Consequently, when a presentation included large tables, the operations proceeded significantly slower. This update provides a patch to optimize the table content traversal. As a result, the operations proceed faster in the described scenario.

BZ#902694

Previously, the keyboard-shortcut mapping was preformed automatically. As a consequence, non-existing keys were suggested as shortcuts in certain languages. With this update, a patch has been provided to fix this bug and affected shortcuts are now mapped manually.

Users of libreoffice are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.96. LIBRTAS

8.96.1. RHEA-2013:1733 — librtas enhancement update

Updated librtas packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

The librtas packages contain a set of libraries that allow access to the Run-Time Abstraction Services (RTAS) on 64-bit PowerPC architectures. The librtasevent library contains definitions and routines for analyzing RTAS events.

Enhancement

BZ#985850

This update adds support for a user space solution for Dynamic Memory Affinity via the PRRN interface. When an affinity for a partition changes as a result of system optimization, the impacted partition will be notified through an event-scan RTAS call that the affinity properties of the partitions have changed. As a result, the partition is expected to refresh its affinity strings through existing RTAS/hidden h_calls.

Users of librtas are advised to upgrade to these updated packages, which add this enhancement.

8.97. LIBTEVENT

8.97.1. RHBA-2013:1552 — libtevent bug fix and enhancement update

Updated libtevent packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The libtevent packages provide Tevent, an event system based on the talloc memory management library. Tevent supports many event types, including timers, signals, and the classic file descriptor events. Tevent also provides helpers to deal with asynchronous code represented by the tevent_req (Tevent Request) functions.



NOTE

The libtevent packages have been upgraded to upstream version 0.9.18, which provides a number of bug fixes and enhancements over the previous version. (BZ#951034)

Bug Fixes

BZ#975489

Prior to this update, a condition in the poll backend copied a 64-bit variable into an unsigned integer variable, which was smaller than 64-bit on 32-bit architectures. Using the unsigned integer variable in a condition rendered the condition to be always false. The variable format has been changed to the uint64_t format guaranteeing its width to be 64 bits on all architectures. As a result, the condition now yields expected results.

BZ#978962

Previously, the tevent_loop_wait() function internally registered its own signal handler even though it had been never removed. Consequently, tevent_loop_wait() could not end even there were no registered custom handlers. This update applies a patch to fix this bug and tevent_loop_wait() now works as expected.

Users of libevent are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.98. LIBVIRT

8.98.1. RHBA-2013:1581 — libvirt bug fix and enhancement update

Updated libvirt packages that fix a number of bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The libvirt library is a C API for managing and interacting with the virtualization capabilities of Linux and other operating systems. In addition, libvirt provides tools for remote management of virtualized systems.

Bug Fixes

BZ#846013

Previously, due to several issues, IPv6 was not handled properly during migration. With this update, migrations now succeed in the described scenario.

BZ#847822

Without manual configuration, the remote driver did not support connection to the session instance of the libvirtd daemon. This behavior could confuse users, who attempted to use such a configuration. With this update, connections that do not have the necessary manual configuration are not allowed by libvirt.

BZ#851075

Previously, the libvirt library was missing driver implementation for the ESX environment. As a consequence, a user could not configure any network for an ESX guest. The network driver has been implemented and a user now can configure networks for ESX guests as expected.

BZ#882077

Previously, libvirt reported raw QEMU errors when creating of snapshots failed, and the error message provided was confusing. With this update, libvirt now gives a clear error message when QEMU is not capable of making snapshots.

BZ#888503

The AMD family 15h processors CPU architecture consists of “modules”, which are represented both as separate cores and separate threads. Management applications needed to choose between one of the approaches, and libvirt did not provide enough information to do this. In addition, the management applications were not able to represent the modules in an AMD family 15h processors core according to their needs. The capabilities XML output now contains more information about the processor topology, so that the management applications can extract the information they need.

BZ#892079

Previously, the libvirtd daemon was unable to execute an s3 or s4 operation for a Microsoft Windows guest which ran the guest agent service. Consequently, this resulted in the “domain s4 fail” error message, due to the domain being destroyed. With this update, the guest is destroyed successfully and libvirtd no longer crashes.

BZ#894723

A virtual machine (VM) can be saved into a compressed file. Previously, when decompression of that

file failed while libvirt was trying to resume the VM, libvirt removed the VM from the list of running VMs. However, it did not remove the corresponding QEMU process. With this update, the QEMU process is killed in such cases. Moreover, non-fatal decompression errors are now ignored and a VM can be successfully resumed if such an error occurs.

BZ#895294

Updating a network interface using the `virDomainUpdateDeviceFlags` API failed when a boot order was set for that interface. The update failed even if the boot order was set in the provided device XML. `virDomainUpdateDeviceFlags` API has been fixed to correctly parse boot order specification from the provided device XML and updating network interfaces with boot orders now works as expected.

BZ#895340

The libvirt library allows users to set Quality of Service (QoS) on a domain's Network Interface Controller (NIC). However, due to a bug in the implementation, certain values were not set correctly. As a consequence, the real throughput did not correspond with the one set in a domain XML. The underlying source code has been modified to set the correct values from the XML and the throughput now corresponds with the one set in the XML as expected.

BZ#895424

Hot unplug of vCPUs is not supported by QEMU in Red Hat Enterprise Linux 6. Therefore, an attempt to use this functionality failed, but the count of processors as remembered by the libvirt library was updated to the new number and remembered. With this update, libvirt now verifies if QEMU actually unplugged the CPUs so that the internal information is updated only when the unplug was successful.

BZ#895826

Previously, when a migration failed, the destination host started to relabel files because it was no longer using them. However, this behavior impacted the source host, which was still running. As a consequence, guests could lose the ability to write to disks. This update applies a patch to fix this bug so that files that are still in use are no longer relabeled in the described scenario.

BZ#895882

Python bindings for the libvirt library contained incorrect implementation of the `getDomain()` and `getConnect()` methods in the `virDomainSnapshot` class. Consequently, the Python client terminated unexpectedly with a segmentation fault. Python bindings now provide the proper `domain()` and `connect()` accessors that fetch Python objects stored internally within the `virDomainSnapshot` instance and crashes no longer occur.

BZ#896013

Previously, the libvirt library added a cache of storage file backing chains, rather than rediscovering the backing chain details on every operation. This cache was then used to decide which files to label for sVirt, but when libvirt switched over to use the cache, the code only populated when the kernel control groups (cgroups) were in use. On setups that did not use cgroups, sVirt was unable to properly label backing chain files due to the lack of backing chain cache information. This behavior caused a regression observed by guests being prevented from running. Now, populating the cache was moved earlier in the process, to be independent of cgroups, the cache results in more efficient sVirt operations, and now works whether or not cgroups are in effect.

BZ#903238

Occasionally, when users ran multiple `virsh create ordestroy` loops, a race condition could occur and the `libvirtd` daemon terminated unexpectedly with a segmentation fault. False error messages regarding the domain having already been destroyed to the caller also occurred. With this update, the

outlined script is run and completes without libvirtd crashing.

BZ#903248

Previously, the libvirt library followed relative backing chains differently than QEMU. This resulted in missing sVirt permissions when libvirt could not follow the chain. With this update, relative backing files are now treated identically in libvirt and QEMU, and VDSM use of relative backing files functions properly.

BZ#903433

When the kernel control group (cgroups) were enabled, moving tasks among cgroups could, in rare occurrences, result in a race condition. Consequently, a guest could fail to start after repeating the start and stop commands tens of times using the virsh utility. With this update, the code that handles groups of threads has been optimized to prevent races while moving from one cgroup to another and guests now start as expected in the described scenario.

BZ#906299

Various memory leaks in the libvirtd daemon were discovered when users ran Coverity and Valgrind leak detection tools. This update addresses these issues, and libvirtd no longer leaks memory in the described scenario.

BZ#908073

Previously, when users started the guest with a sharable block CD-Rom, the libvirtd daemon failed unexpectedly due to accessing memory that had been already freed. This update addresses the aforementioned issue, and libvirtd no longer crashes in the described scenario.

BZ#911609

Due to a race condition in the libvirt client library, any application using libvirt could terminate unexpectedly with a segmentation fault. This happened when one thread executed the connection close callback, while another one freed the connection object, and the connection callback thread then accessed memory that had been already freed. This update fixes the possibility of freeing the callback data when they are still being accessed.

BZ#912179

When asked to create a logical volume with zero allocation, the libvirt library ran the lvcreate command to create a volume with no extends, which is not permitted. Creation of logical volumes with zero allocation failed and libvirt returned an error message that did not mention the correct error. Now, rather than asking for no extends, libvirt tries to create the volume with a minimal number of extends. The code has been also fixed to provide the correct error message when the volume creation process fails. As a result, logical volumes with zero allocation can now be successfully created using libvirt.

BZ#913244

When auto-port and port were not specified, but the tlsPort attribute was set to "-1", the tlsPort parameter specified in the QEMU command line was set to "1" instead of a valid port. Consequently, QEMU failed, because it was unable to bind a socket on the port. This update replaces the current QEMU driver code for managing port reservations with the new virPortAllocator APIs, and QEMU is now able to bind a socket on the port.

BZ#913363

The libvirt library could abort migration when domain's disks used unsafe cache settings even though they were not stored on a shared storage and libvirt was explicitly asked to copy all storage. As a

consequence, migration without a shared storage was only possible with the VIR_MIGRATE_UNSAFE flag enabled. With this update, the test for safe disk cache settings is now limited only to shared storage because any setting is safe for locally stored disk images.

BZ#914677

Previously, the libvirt library was not tolerant of missing unpriv_sgio support in running kernel even though it was not necessary. Consequently, after upgrading the host system to Red Hat Enterprise Linux 6.5, users were unable to start domains using shareable block disk devices unless they rebooted the host into the new kernel. With this update, the check for unpriv_sgio support is only performed when it is really needed. As a result, libvirt is now able to start all domains that do not strictly require unpriv_sgio support regardless of host kernel support for it.

BZ#916315

Due to a bug in the libvirt code, two APIs, vidDomainBlockStatsFlags() and vidDomainDetachDeviceFlags(), were executed concurrently. As a consequence, the libvirtd daemon terminated unexpectedly. The underlying source code has been modified to make these APIs mutually exclusive so that the daemon no longer crashes in such a case.

BZ#917510

When a virtual machine (VM) with a managed save image was started with the "--force-boot" parameter that removed the managed save image, a flag holding the managed save state was not cleared. This caused that incorrect information was displayed and some operations regarding managed save state failed. This bug has been fixed and the flag is now correctly cleared in the described scenario.

BZ#920205

At the end of migration, libvirt was waiting for the Simple Protocol For Computing Environments (SPICE) data to be migrated to the destination QEMU, before it resumed the domain on the destination host. This significantly increased the waiting time when the domain was not running on any host. With this update, the underlying code has been modified to not to wait until the end of the SPICE migration. As a result, the resume is done as soon as possible without any significant delay.

BZ#920441

Previously, the listen attribute in QEMU cookie files was discarded. Consequently, if the user had different networks in use, one for management and migration, and one for Virtual Network Computing (VNC) and SPICE, the remote host name was passed to QEMU via the client_migrate_info flag. This caused the SPICE client to be disconnected upon migration of a virtual machine. With this update, the remote listen address is passed instead and the SPICE client is no longer disconnected in the described scenario.

BZ#921387

Due to the use-after-free bug in the logical storage back end, the libvirtd daemon could terminate unexpectedly when deleting the logical storage pool. The underlying source code has been modified and the daemon now works as expected when deleting logical volumes.

BZ#921538

Due to a race condition in the client side of libvirt's RPC implementation, a client connection that was closed by the server could be freed, even though other threads were still waiting for APIs sent through this connection to finish. As a consequence, the other threads could have accessed memory that had already been freed and the client terminated unexpectedly with a segmentation fault. With this update the connection is freed only after all threads process their API calls and report errors to their callers.

BZ#921777

Previously, a lock used when dealing with transient networks was incorrect. Consequently, when the `define` API was used on a transient network, the network object lock was not unlocked as expected. The underlying source code has been modified and the object lock is now unlocked correctly.

BZ#922153

Previously, the `libvirt` library made control group (cgroup) requests on files that it should not have. With older kernels, such nonsensical cgroup requests were ignored; however, newer kernels are stricter, resulting in `libvirt` logging spurious warnings and failures to the `libvirtd` and audit logs. The audit log failures displayed by the `ausearch` tool were similar to the following:

```
root [date] - failed cgroup allow path rw /dev/kqemu
```

With this update, `libvirt` no longer attempts the nonsensical cgroup actions, leaving only valid attempts in the `libvirtd` and audit logs.

BZ#922203

Previously, the `libvirt` library used the incorrect variable when constructing audit messages. This led to invalid audit messages, causing the `ausearch` utility to format certain entries as having “`path=(null)`” instead of the correct path. This could prevent `ausearch` from locating events related to cgroup device Access Control Lists (ACL) modifications for guests managed by `libvirt`. With this update, the audit messages are generated correctly, preventing loss of audit coverage.

BZ#923613

Previously, the `vol-download` command was described incorrectly in the `virsh(1)` manual page. With this update, the command description has been fixed.

BZ#923946

When SELinux was disabled on a host, or the QEMU driver was configured not to use it, and the domain XML configuration contained an explicit `seclabel` option, the code parsed the `seclabel` option, but ignored it later when it was generating labels on domain start, and created a new and empty `seclabel` entry [`seclabeltype='none'`]. Consequently, a migration between two hosts running Red Hat Enterprise Linux 6.5 failed with the following error message:

```
libvirtError: XML error: missing security model when using multiple labels
```

With this update, if the `seclabel` entry already exists, a new one is no longer created, and the migration works as expected in the described scenario.

BZ#923963

Previously, there was an Application Binary Interface (ABI) inconsistency in messages of the kernel netlink protocol between certain versions of Red Hat Enterprise Linux. When the `libvirt` library sent a netlink `NLM_F_REQUEST` message and the `libvirt` binary had been built using kernel header files from a different version of the kernel than the version of the machine running `libvirt`, errors were returned. Consequently, Peripheral Component Interconnect (PCI) passthrough device assignments of SR-IOV network devices failed when they used the [`interface type='hostdev'`] option, or when the `libvirt` network was set with the [`forward mode='hostdev'`] option. In such a case, the following error message or a similar one was returned:

```
error dumping (eth3) (3) interface: Invalid argument
```

With this update, libvirt retries the NLM_F_REQUEST message formatted appropriately for all versions of the kernel. Now, a single libvirt binary successfully assigns SR-IOV network devices to a guest using PCI passthrough on a host running any version of Red Hat Enterprise Linux 6 kernel.

BZ#924571

Previously, the vol-name command of the virsh utility printed a NULL string when there was no option for specifying the pool. Consequently, an error message was returned, which could confuse users. The command has been modified to not require to specify an option in case where it is not needed. As a result, the error message is no longer returned in the described scenario.

BZ#924648

The QEMU driver currently does not support increasing of the maximum memory size. However, this ability was documented in the virsh(1) manual page. With this update, the manual page has been corrected.

BZ#928661

Previously, part of the code refactoring to fix another bug, left a case where locks were cleaned up incorrectly. As a consequence, the libvirtd daemon could terminate unexpectedly on certain migration to file scenarios. After this update, the lock cleanup paths were fixed and libvirtd no longer crashed when saving a domain to a file.

BZ#947387

The libvirt library uses side files to store the internal state of managed domains in order to re-read the state upon the libvirtd service restart. However, if a domain state was saved in an inconsistent state, the state was not re-read and the corresponding domain was lost. As a consequence, the domain could disappear. After this update, when the libvirtd service is saving the internal state of a domain, the consistent internal state is saved and domains which may break it are disallowed from starting. As a result, the domain is no longer forgotten.

BZ#948678

Previously, attempts to clone a storage volume that was not in the RAW format from a directory pool, file system pool, or NFS pool, to a LVM pool, using the “virsh vol-create-from” command, failed with an “unknown file format” error message. This update fixes this bug by treating output block devices as the RAW file format and storage volumes can now be cloned as expected.

BZ#950286

Under certain conditions, when a connection was closed, guests set to be automatically destroyed failed to be destroyed. As a consequence, the libvirtd daemon terminated unexpectedly. A series of patches addressing various crash scenarios has been provided and libvirtd no longer crashes while auto-destroying guests.

BZ#951227

When running the libvirt test suite on a machine under a heavy load, the test could end up in a deadlock. Since the test suite was run during an RPM build, the build never finished if a deadlock occurred. This update fixes the handling of an event loop used in the test suite, and the test suite no longer hangs in the described scenario.

BZ#955575

Previously, the VirtualHW application version 9 was not set as supported even though the corresponding ESX version 5.1 was set to be supported earlier. As a consequence, when a connection was made to an ESX 5.1 server with a guest using virtualHW version 9, the following error

was displayed:

```
internal error Expecting VMX entry 'virtualHW.version' to be 4, 7 or 8
but found 9
```

This update adds VirtualHW version 9 into the list of supported versions and the aforementioned error message is no longer displayed in this scenario.

BZ#960683

Libvirt's internal data structures which hold information about the topology of the host and guest, are limited in size to avoid the possibility of a denial-of-service (DoS) attack on the daemon. However, these limits were too strict and did not take into account the possibility that hosts with 4096 CPUs might be used with libvirt. After this update, the limits have been increased to allow scalability even on larger systems.

BZ#961034

Prior to this update, the `F_DUPFD_CLOEXEC` operation with the `fcntl()` function expected a single argument, specifying the minimum file descriptor (FD) number, but none was provided. Consequently, random stack data were accessed as the FD number and a libvirt live migration could then terminate unexpectedly. This update ensures that the argument is provided in the described scenario, thus fixing this bug.

BZ#964359

Previously, the `libvirtd` daemon set up supplemental groups of child processes by making a call between the `fork()` and `exec()` functions to the `getpwuid_r()` function, which could cause a mutual exclusion (mutex). As a consequence, if another thread was already holding the `getpwuid_r` mutex at the time `libvirtd` called the `fork()` function, the forked child process deadlocked, which in turn caused `libvirtd` to become unresponsive. The code to compute the set of supplemental groups has been refactored so that no mutex is required after `fork`. As a result, the deadlock scenario is no longer possible.

BZ#965442

Previously, the libvirt library did not update the pool information after adding, removing, or resizing a volume. As a consequence, the user had to refresh the pool using the "virsh pool-refresh" command to get the correct pool information after these actions. After this update, the pool information is automatically updated after adding, removing, or resizing a volume.

BZ#970495

Previously, the `virsh` utility considered the "--pool" argument of the "vol-create" and "vol-create-as" commands to be a pool name. As a consequence, `vol-create` and `vol-create-as` `virsh` commands did not work when a pool was specified by its Universally Unique Identifier (UUID), even though they were documented to accept both name and UUID for pool specification. With this update, `virsh` has been fixed to look up a pool both by name and UUID. As a result, both `virsh` commands now work according to their documentation.

BZ#971485

Previously, if the user had not specified a Virtual Network Computing (VNC) address in their domain XML, the one from the `qemu.conf` file was used. However, upon migrating, there was no difference between cases where the listen address was set by user in the XML directly or copied from the `qemu.conf` file. As a consequence, a domain could not be migrated. After this update, if the listen address is copied from `qemu.conf`, it is not transferred to the destination. As a result, a domain can be migrated successfully.

BZ#971904

Previously, the libvirt library's logging function that was passed to the libudev library did not handle strings with multiple parameters correctly. As a consequence, the libvirtd daemon could terminate unexpectedly when libudev logged a message. After this update, libvirt now handles multiple parameters correctly. As a result, libvirtd no longer crashes when libudev logs messages.

BZ#975201

Previously, the libvirt library only loaded one Certification Authority (CA) certificate from the cacert.pem file even though the file contained several chained CA certificates. As a consequence, libvirt failed to validate client and server certificates when they were both signed by intermediate CA certificates, sharing a common ancestor CA. After this update, the underlying code has been fixed to load all CA certificates. As a result, the CA certificate validation code correctly works when a client and server certificates are both signed by intermediate CA certificate, sharing a common ancestor CA.

BZ#975751

Previously, due to loader Hypervisor versions, many features were available only for guests with only one display. As a consequence, guests with two displays could not properly be defined on the QEMU hypervisor and some other features were not properly taking the second display into consideration. With this update, the ability to define more display types and all one-display assumptions were fixed in all relevant code. As a result, domains with multiple displays can now be defined, properly migrated, and started.

BZ#976401

The SPICE protocol can be set to listen on the given IP address or obtain the listening IP address from the given network. QEMU does not allow changing the SPICE listening IP address at runtime, therefore the libvirt library verifies this IP address with every user's update of SPICE settings on a guest. A regression bug in the libvirt code caused libvirt to incorrectly evaluate this listening IP address check if the user had SPICE set to listen on the given network because the user's XML request contained both, the listening IP address and network address. Consequently, the user's operation was rejected. With this update, libvirt considers also the type of the listening IP address when comparing an IP address from the user's request with the current listening IP address. The user is now able to update SPICE settings on a guest as expected in this scenario.

BZ#977961

When migrating, the libvirtd daemon leaked migration Uniform Resource Identifier (URI) on a destination guest. A patch has been provided to fix this bug and the migration URI is now freed correctly.

BZ#978352

Prior to this update, the libvirtd daemon leaked memory in the virCgroupMoveTask() function. A fix has been provided which prevents libvirtd from incorrect management of memory allocations.

BZ#978356

Previously, the libvirtd daemon was accessing one byte before the array in the virCgroupGetValueStr() function. This bug has been fixed and libvirtd now stays within array bounds.

BZ#979330

Previously, the libvirt library depended on a "change" notification from the kernel to indicate that it should change the name of the device driver bound to a device. However, this change notification was not sent. As a consequence, the output from the "virsh nodedev-dumpxml" command always

showed the device driver that was bound to the device at the time libvirt was started and not the currently-bound driver. This bug has been fixed and libvirt now manually updates the driver name every time a "nodedev-dumpxml" command is executed, rather than depending on a change notification. As a result, the driver name from the output of "nodedev-dumpxml" is always correct.

BZ#980339

Previously, if an incorrect device name was given in the <pf> element of a libvirt network definition, libvirt terminated unexpectedly when a guest attempted to create an interface using that network. With this update, libvirt now validates the <pf> device name to verify that it exists and that it is an sriov-capable network device. As a result, libvirt no longer crashes when a network with incorrect <pf> is referenced. Instead, it logs an appropriate error message and prevents the operation.

BZ#983539

Previously, the virStorageBackendFileSystemMount() function returned success even if the mount command had failed. As a consequence, libvirt showed the pool as running even though it was unusable. After this update, an error is displayed if the mount command has failed. As a result, libvirt no longer displays a success message when the mount command fails.

BZ#999107

Due to an omission in the libvirt code, the VLAN tag for a hostdev-based network (a network which is a pool of SRIOV virtual functions to be assigned to guests via PCI device assignment) was not being properly set in the hardware device. With this update, the missing code has been provided and a VLAN tag set in the network definition is now properly presented to the devices as they are assigned to guests.

BZ#1001881

Previously, the libvirt library was erroneously attempting to use the same alias name for multiple hostdev network devices. As a consequence, it was impossible to start a guest that had more than one hostdev network device in its configuration. With this update, libvirt now ensures that each device has a different alias name. As a result, it is now possible to start a guest with multiple hostdev network devices in its configuration.

BZ#1002790

The description of the blockcopy command in the virsh(1) manual page was identical to the description of the blockpull command. The correct descriptions have been provided with this update.

BZ#1006710

Previously, when parsing the domain XML with an "auto" numatune placement and the "nodeset" option was specified, the nodeset bitmap was freed twice. As a consequence, the libvirtd daemon terminated unexpectedly due to the double freeing. After this update, libvirtd now sets the pointer to NULL after freeing it. As a result, libvirtd no longer crashes in this scenario.

BZ#1009886

Previously, due to code movement, there was an invalid job used for querying for the SPICE migration status. As a consequence, when migrating a domain with a Simple Protocol for Independent Computing Environments (SPICE) seamless migration and using the domjobinfo command to request information on the same domain at the same time, the libvirtd daemon terminated unexpectedly. After this update, the job has been set properly and libvirtd no longer crashes in this scenario.

BZ#1011981

Whereas the status command of libvirt-guests init script returned the "0" value when libvirt-guests

service was stopped, Linux Standard Base (LSB) required a different value (“3”) in such case. Consequently, other scripts relying on the return value could not distinguish whether the service was running or not. The libvirt-guests script has been fixed to conform with LSB and the “service libvirt-guests status” command now returns the correct value in the described scenario.

BZ#1013758

Previously, the libvirt library contained a heuristic to determine the limit for maximum memory usage by a QEMU process. If the limit was reached, the kernel just killed the QEMU process and the domain was killed as well. This, however, cannot be guessed correctly. As a consequence, domains were killed randomly. With this update, the heuristic has been dropped and domains are not killed by the kernel anymore.

Enhancements**BZ#803602**

This enhancement adds the ability to specify a share policy for domain's Virtual Network Computing (VNC) console. Latest changes in QEMU behavior from shared to exclusive VNC caused certain deployments, which used only shared VPN, to stop working. With a new attribute, “sharePolicy”, users are able to change the policy from exclusive to share and such deployments now work correctly.

BZ#849796

This enhancement introduces QEMU's native GlusterFS support. Users are now able to add a disk image stored on the GlusterFS volumes to a QEMU domain as a network disk.

BZ#851455

Due to security reasons, the libvirt library uses by default only ports larger than 1023 (“unprivileged ports”) for Network Address Translation (NAT) of network traffic from guests. However, sometimes the guests need access to network services that are only available if a privileged port is used. This enhancement provides a new element, “<nat>”, which allows the user to specify both a port or an address range to use for NAT of network traffic.

BZ#878765

This update adds a missing description about the “migrateuri” parameter of the “migrate” command to the virsh(1) manual page.

BZ#896604

With this enhancement, the libvirt library now supports the ram_size parameter. Users are now able to set the RAM memory when using multiple heads in one Peripheral Component Interconnect (PCI) device.

BZ#924400

The QEMU guest agent now supports enabling and disabling of guest CPUs. With this enhancement, support for this feature has been added to the libvirt library so that users are now able to use libvirt APIs to disable CPUs in a guest for performance and scalability reasons.

BZ#928638

Domain Name System (DNS) servers and especially root DNS servers, discourage forwarding of DNS requests that are not fully qualified domain names, that is, which include the domain as well as the host name. Also, the dnsmasq processes started by libvirt to service guests on its virtual networks

prohibit forwarding such requests. However, there are certain circumstances where this is desirable. This update adds the permission for upstream forwarding of (DNS) requests with unqualified domain names. The libvirt library now provides an option in its network configuration to allow forwarding of DNS requests with non-qualified hostnames. The "forwardPlainNames='yes'" option must be added as an attribute to the <dns> element of a network, after which such forwards are allowed.

BZ#947118

Support for locking a domain's memory in the host's memory has been added to the libvirt library. This update enables users to avoid domain's memory pages to be swapped, and thus to avoid the latency in domain execution caused by swapping. Users can now configure domains to always be present in the host memory.

BZ#956826

QEMU I/O throttling provides a fine-grained I/O control in virtual machines and provides an abstraction layer on top of the underlying storage devices.

BZ#826315, BZ#822306

A new pvpanic virtual device can be wired into the virtualization stack and a guest panic can cause libvirt to send a notification event to management applications. This feature is introduced in Red Hat Enterprise Linux 6.5 as a Technology Preview. Note that enabling the use of this device requires the use of additional qemu command line options; this release does not include any supported way for libvirt to set those options.

BZ#1014198

Previously, the virDomainDeviceUpdateFlags() function in the libvirt library allowed users to update some configuration on a domain device while the domain was still running. Consequently, when updating Network Interface Controller (NIC), the QoS could not be changed because of a missing implementation. With this update, the missing implementation has been added, and QoS can now be updated on a NIC.

Users of libvirt are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. After installing the updated packages, libvirtd will be restarted automatically.

8.98.2. RHBA-2013:1748 — libvirt bug fix update

Updated libvirt packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The libvirt library is a C API for managing and interacting with the virtualization capabilities of Linux and other operating systems. In addition, libvirt provides tools for remote management of virtualized systems.

Bug Fix

BZ#1029632

When two clients tried to start the same transient domain, libvirt may have not properly detected that the same domain had already been being started. Consequently, more than one QEMU process could run for the same domain while libvirt did not know about them. With this update, libvirt has been fixed to properly check whether the same domain is not already being started, and thus avoids starting more than one QEMU process for the same domain.

Users of libvirt are advised to upgrade to these updated packages, which fix this bug. After installing the updated packages, libvirtd will be restarted automatically.

8.99. LIBVIRT-CIM

8.99.1. RHBA-2013:1676 — libvirt-cim bug fix update

Updated libvirt-cim packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The libvirt-cim packages contain a Common Information Model (CIM) provider based on Common Manageability Programming Interface (CMPI). It supports most libvirt virtualization features and allows management of multiple libvirt-based platforms.

Bug Fixes

BZ#826179

Previously, running the wbemcli utility with the KVM_ComputerSystem class terminated unexpectedly with a segmentation fault. This was because even when connecting to the libvirtd daemon read-only, the domain XML with secure information, that is with the VIR_DOMAIN_XML_SECURE flag, was dumped. However, this operation is forbidden in libvirt. With this update, the flag is not used with read-only connections. Running the wbemcli command with KVM_ComputerSystem now displays the domain information as expected.

BZ#833633

When updating certain libvirt-cim or sblim-smis-hba packages, the following error could have been logged in the /var/log/messages file:

```
sfcbmof: *** Repository error for /var/lib/sfcb/registration/repository//root/pg_interop/qualifiers
```

This problem occurred because libvirt-cim installed the PG_InterOp class incorrectly in the sblim-sfcb repository, however, this class is specific for the open-pegasus package. With this update, PG_InterOp is unregistered before upgrading the package, and no error message is logged in this scenario.

BZ#859122

Previously, libvirt-cim incorrectly installed providers specific for the open-pegasus package in the sblim-sfcb repository. This could have caused various problems, for example, failures when compiling the MOF files. Providers specific for open-pegasus are now installed in the correct repository and the problems no longer occur.

BZ#908083

Previously, if a qemu domain was defined with a bridge network interface, running the libvirt-cim provider failed with the following error message:

```
Unable to start domain: unsupported configuration: scripts are are not supported on interfaces of type bridge
```

This was because code triggering a script was added in a file used to create the domain prior to checking the qemu domain type. However, scripts are not allowed for qemu domains. With this update, a check for the qemu domain type is performed prior to adding the code triggering the script. As a result, when using libvirt-cim, it is now possible to create qemu domains with the bridge network interface.

BZ#913164

Previously, a call to query a guest's current VNC address and port number returned the static configuration of the guest. If the guest was used to enable the "autoport" selection, the call did not

return the allocated port. The libvirt-cim code has been modified to only return static configuration information. This allows other interfaces to return information based on the domain state. As a result, the current and correct port being used by the domain for VNC is now returned.

BZ#1000937

Virtual machines managed by a libvirt-cim broker were not aware of the "dumpCore" flag in the "memory" section nor was there support for the "shareable" property for "disk" devices. Thus, those properties were dropped from the virtual machine XML configuration when the configuration was updated by the broker. As a consequence, customers expecting or setting these properties on their virtual machines had to adjust the configurations in order to reset them. With this update, a patch has been added to libvirt-cim and it is now aware of these properties so that no changes made to the virtual machine XML configuration will be lost by the broker when it writes the configuration. As a result, virtual machines managed by the libvirt-cim broker will recognize the "dumpCore" tag in the "memory" section or the "shareable" tag on a "disk" device and not remove either when updating the virtual machine XML configuration.

Users of libvirt-cim are advised to upgrade to these updated packages, which fix these bugs.

8.100. LIBVIRT-SNMP

8.100.1. RHBA-2013:1666 — libvirt-snmpp bug fix update

Updated libvirt-snmpp packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The libvirt-snmpp packages allow users to control and monitor the libvirt virtualization management tool through Simple Network Management Protocol (SNMP).

Bug Fix

BZ#736258

Previously, closing the libvirtMib_subagent using the Ctrl+C key combination led to a memory leak. The libvirtd daemon could be also terminated sometimes. A patch has been applied to address this issue, and a memory leak no longer occurs in this scenario.

Users of libvirt-snmpp are advised to upgrade to these updated packages, which fix this bug.

8.101. LIBWACOM

8.101.1. RHBA-2013:1567 — libwacom bug fix update

Updated libwacom packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The libwacom packages contain a library that provides access to a tablet model database. The libwacom packages expose the contents of this database to applications, allowing for tablet-specific user interfaces. The libwacom packages allow the GNOME tools to automatically configure screen mappings and calibrations, and provide device-specific configurations.

Bug Fix

BZ#847427

Previously, the Wacom Stylus pen was not supported on Lenovo ThinkPad X220 tablets by the

libwacom database. Consequently, the pen was not recognized by the `gnome-wacom-properties` tool, and warning messages were returned. Support for the Wacom Stylus on Lenovo ThinkPad X220 tablets has been added and `gnome-wacom-properties` is now able to calibrate the tablet.

Users of `libwacom` are advised to upgrade to these updated packages, which fix this bug.

8.102. LIBXML2

8.102.1. RHBA-2013:1737 — libxml2 bug fix update

Updated `libxml2` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `libxml2` library is a development toolbox providing the implementation of various XML standards.

Bug Fix

BZ#863166

Previously, parsing an XML file containing entities loaded via Document Type Definition (DTD) using the `XML::LibXML` module could lead to a missing entity error as `XML::LibXML` did not load entities DTD. A patch has been applied to address this problem and XML files are parsed successfully in this scenario.

Users of `libxml2` are advised to upgrade to these updated packages, which fix this bug. The desktop must be restarted (log out, then log back in) for this update to take effect.

8.103. LINUXPTP

8.103.1. RHBA-2013:1564 — linuxptp bug fix and enhancement update

Updated `linuxptp` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Linux PTP project is a software implementation of the Precision Time Protocol (PTP) according to IEEE standard 1588 for Linux. These packages provide a robust implementation of the standard and use the most relevant and modern Application Programming Interfaces (API) offered by the Linux kernel. Supporting legacy APIs and other platforms is not a goal.



NOTE

The `linuxptp` package has been upgraded to upstream version 1.3, which provides a number of bug fixes and enhancements over the previous version. (BZ#916787)

Bug Fixes

BZ#910966

Previously, the `ptp4l` application did not limit the frequency correction of the clock. As a consequence, with some PTP clocks, when `ptp4l` was correcting a large offset, it could set the frequency correction to -100%, which effectively stopped the clock. This update adds a new option to configure the maximum allowed correction of the clock, which, by default is 90%. As a result, the synchronized clock never stops unless `ptp4l` is allowed to adjust the clock by 100%.

BZ#910974

Previously, the phc2sys utility was not able to read information about the current Coordinated Universal Time (UTC) offset and pending leap seconds from the ptp4l application. As a consequence, the user had to specify the UTC offset manually and the leap seconds were not handled. This update adds a new option to phc2sys to wait for ptp4l to synchronize the PTP clock and to periodically read the current UTC offset and information about pending leap seconds. As a result, the phc2sys utility uses the correct UTC offset and leap seconds are handled properly.

BZ#991332, BZ#985531

Previously, the ptp4l application did not correctly check if a cached follow-up or a synchronized message could be associated with a newly received synchronization or a follow-up message. As a consequence, the messages could be associated incorrectly, which could result in a large offset and disturbed synchronization of the clock. The code which associates the synchronization and follow-up messages has been fixed. As a result, there are no longer disturbances in the synchronization.

BZ#991337

Previously, the ptp4l application did not reset the announce receipt timer for ports in the PASSIVE state when an announce message was received. As a consequence, the port in the PASSIVE state was repeatedly switching between PASSIVE and MASTER states. This bug has been fixed and the timer is now correctly reset with every announce message. As a result, the port stays in the PASSIVE state until it stops receiving announce messages.

BZ#966787

Previously, the ptp4l and phc2sys utilities did not check if the command line arguments and the values specified in the configuration file were valid. As a consequence, the utilities could terminate unexpectedly. The utilities now check if the values are valid and if an invalid value is specified, the utilities no longer terminate unexpectedly and print an error message instead.

Enhancement**BZ#977258**

Occasionally, it is important that the system clock is not stepped, that is, not to interfere with other programs running on the system. Restarting the phc2sys application caused stepping of the clock. A new option has been added to phc2sys, and it is now possible to prevent phc2sys from stepping the clock.

Users of linuxptp are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.104. LKSCTP-TOOLS**8.104.1. RHBA-2013:1726 — lksctp-tools bug fix and enhancement update**

Updated lksctp-tools packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

These packages are intended to supplement the Stream Control Transmission Protocol (SCTP) implementation, which has been a part of the kernel since kernel version 2.5.36. For more information on LKSCTP see the section titled "LKSCTP - Linux Kernel SCTP" in the README file included in the package documentation. These packages contain the base runtime library and command line tools.

**NOTE**

The lksctp-tools packages have been upgraded to upstream version 1.0.10, which provide a number of bug fixes and enhancements over the previous version. The patches include updates in the header file which enable users to make use of new SCTP kernel features, for example, the introduction of `SCTP_GET_ASSOC_STATS` socket option in order to retrieve association statistics. (BZ#[855379](#), BZ#[908390](#), BZ#[912557](#), BZ#[953383](#))

Users of lksctp-tools are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.105. LOGROTATE

8.105.1. RHBA-2013:1095 — logrotate bug fix update

Updated logrotate packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The logrotate utility simplifies the administration of multiple log files, allowing the automatic rotation, compression, removal, and mailing of log files.

Bug Fixes

BZ#[841520](#)

The logrotate utility always tried to set owner of the rotated log even when the owner was the same as the current owner of the log file. Consequently, the rotation failed on file systems or systems where changing the ownership was not supported. With this update, before the ownership is changed, logrotate check if it is a real ownership change; that is, logrotate verifies if the new ownership is not the same as the previous one, and skips the change if the ownership change has not been real. The logrotate utility now rotates logs as expected in this scenario.

BZ#[847338](#)

Setting the Access control list (ACL) on a rotated log overwrote the previously set mode of the log file. As a consequence, the "create" directive was ignored. To fix this bug, the ACL is no longer copied from the old log file when using the "create" directive and the mode defined using the "create" directive is used instead. As a result, "create" mode works as expected and it is no longer ignored in the described scenario.

BZ#[847339](#)

Both the `acl_set_fd()` and `fchmod()` functions were called to set the log files permissions. Consequently, there was a race condition where the log file could have unsafe permissions for a short time during its creation. With this update, only one of those functions is now called depending on directives combination used in the configuration file and race condition between the `acl_set_fd()` and `fchmod()` function is not possible in the described scenario.

BZ#[848131](#)

Because the inverse umask value 0000 was used when creating a new log file, the newly created log file could have unwanted 0600 permissions for a short time before the permissions were set to the proper value using the `fchmod()` function. With this update, umask is set to 0777 and the newly created log file has proper 0000 permissions for this short period.

BZ#[920030](#)

The default SELinux context was set after the compressed log file had been created. Consequently, the compressed log did not have the proper SELinux context. With this update, the default SELinux context is now set before the compressed log file creation and compressed log files have proper SELinux context.

BZ#922169

Temporary files created by the logrotate utility were not removed if an error occurred during its use. With this update, temporary files are now removed in such a case.

Users of logrotate are advised to upgrade to these updated packages, which fix these bugs.

8.106. LOGWATCH

8.106.1. RHBA-2013:1247 — logwatch bug fix update

An updated logwatch package that fixes several bugs is now available for Red Hat Enterprise Linux 6.

Logwatch is a customizable, pluggable log-monitoring system. It will go through the user's logs for a given period of time and make a report in the areas that the user needs.

Bug Fixes

BZ#737247

Previously, logwatch did not correctly parse the up2date service's "updateLoginInfo() login info" messages and displayed them as unmatched entries. With this update, parsing of such log messages has been fixed and works as expected.

BZ#799690

Prior to this update, logwatch did not correctly parse many Openswan log messages and displayed them as unmatched entries. With this update, parsing of such log messages has been fixed and works as expected.

BZ#799987

Logwatch did not parse Dovecot 2.x log messages properly. That resulted in a lot of unmatched entries in its reports. This patch adds additional logic to correctly parse Dovecot 2.x logs, thus unmatched entries related to Dovecot 2.x messages no longer appear.

BZ#800843

The .hdr files are headers for RPM packages; they are essentially metadata. Logwatch's HTTP service parser emitted warnings for the .hdr files, even when the "Detail" parameter was set to "Low". With this update, the .hdr files are now parsed as archives, which removes spurious warnings about the .hdr files.

BZ#837034

Previously, logwatch did not correctly handle the "MailTo" option in its configuration. That resulted in no output, even though a report should have been displayed. This patch adds additional logic to correctly handle an empty "MailTo" option. As a result, output is correctly produced even when this option is empty.

BZ#888007

Prior to this update, logwatch did not correctly parse many smartd log messages and displayed them as unmatched entries. With this update, parsing of such log messages has been fixed and works as expected.

BZ#894134

Prior to this update, logwatch did not correctly parse DNS log messages with DNSSEC validation enabled and displayed them as unmatched entries. With this update, parsing of such log messages has been fixed and works as expected.

BZ#894185

Previously, logwatch did not correctly parse the postfix service's "improper command pipelining" messages and displayed them as unmatched entries. With this update, parsing of such log messages has been fixed and works as expected.

BZ#894191

Previously, logwatch did not correctly parse user names in the secure log. It improperly assumed that such names are composed of letters only and displayed messages containing names with other symbols, such as digits, as unmatched entries. With this update, parsing of user names has been enhanced to include underscores and digits, thus log messages containing such user names no longer display as unmatched entries.

BZ#974042

Logins initiated with the "su -" or "su -l" command were not correctly parsed by logwatch and were displayed as unmatched entries. This update fixes this bug.

BZ#974044

Prior to this update, logwatch did not correctly parse the RSYSLOG_FileFormat time stamps and displayed them as unmatched entries. With this update, parsing of the rsyslog time stamps has been fixed and works as expected.

BZ#974046

SSH Kerberos (GSS) logins were not correctly parsed by logwatch and were displayed as unmatched entries. This update fixes this bug.

BZ#974047

Xen virtual console logins were not correctly parsed by logwatch and were displayed as unmatched entries. This update fixes this bug.

Users of logwatch are advised to upgrade to this updated package, which fixes these bugs.

8.107. LUCI

8.107.1. [RHSA-2013:1603](#) — Moderate: luci security, bug fix, and enhancement update

Updated luci packages that fix two security issues, several bugs, and add two enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

Luci is a web-based high availability administration application.

Security Fixes

CVE-2013-4482

A flaw was found in the way the luci service was initialized. If a system administrator started the luci service from a directory that was writable to by a local user, that user could use this flaw to execute arbitrary code as the root or luci user.

CVE-2013-4481

A flaw was found in the way luci generated its configuration file. The file was created as world readable for a short period of time, allowing a local user to gain access to the authentication secrets stored in the configuration file.

These issues were discovered by Jan Pokorný of Red Hat.

Bug Fixes

BZ#917747

Previously, luci did not reflect concurrent additions to fence devices coverage as happened in the fence-agents package. Consequently, Dell iDRAC (idrac), HP iLO2 (ilo2), HP iLO3 (ilo3), HP iLO4 (ilo4), and IBM Integrated Management Module (imm) devices or agents were not honored in luci, leading to the inability to properly work with or to setup a cluster comprising of these devices. This update restores the capability of luci to work with a full intended set of fence devices.

BZ#956360

Previously, luci did not run in FIPS mode because it utilized components that were not compliant with FIPS. Both components, the python-breaker library and the python-weberror error handler have been modified to comply with FIPS so that luci now works in FIPS mode as expected.

BZ#978479

Due to a bug in the luci code, a data race condition could occur while adding multiple nodes into a cluster with a single request. As a consequence, nodes could have been provided configurations with varying version numbers, leaving the cluster in an unexpected state. The respective luci code has been fixed so this data race cannot be triggered anymore. Multiple nodes can now be added to a cluster at once without a risk of negative consequences.

BZ#773491

Previous implementation of dynamic pop-up messages had a high probability of messages leaving the screen unnoticed under certain circumstances. Therefore, the respective luci code has been modified to adjust dynamic pop-ups to appear as static messages, which significantly decreases a chance that the message might be unnoticed.

BZ#883008

Previously, luci did not reflect concurrent additions to parameters for some fence devices (including "cmd_prompt", "login_timeout", "power_timeout", "retry_on", "shell_timeout") or respective instances ("delay") as happened in the fence-agents package. Consequently, the valid parameters could be

dropped from the respective part of the configuration upon submitting the dedicated forms in luci. This update restores the capability of luci to work with a full intended set of fence agents parameters and, in turn, prevents luci from unexpectedly discarding the already configured parameters.

BZ#896244

Due to a bug in the `cluster.conf(5)` man page, luci expected the default value for the `syslog_facility` option in the cluster logging configuration to be "daemon" instead of the actual default value "local4". Consequently, all logging configuration items without "syslog_facility" explicitly set were thus marked as having "Syslog Message Facility" of "daemon" in luci. This could result in no cluster messages being logged into the custom log file for the rules containing "daemon.*". With this update, luci correctly recognizes "local4" as the default syslog message facility and logging configuration items in luci are marked accordingly by default. The user is now able to effectively set the syslog facility of the logging configuration item to be "daemon". In such a case, cluster messages are logged into log files containing the "daemon.*" rules as expected.

BZ#886517

The luci application did not automatically enable the `ricci` and `modclusterd` services upon creating a new cluster or adding a node to the existing cluster. Therefore, an administrator's intervention was necessary because these services are essential for managing the cluster during its life-cycle. Without these services, luci sustained the contact with cluster nodes, preventing the cluster from rebooting. With this update, luci has been modified to enable the `ricci` and `modclusterd` services on every cluster's node when creating a new cluster or adding a node to the existing cluster. The administrator's intervention is no longer needed in the aforementioned scenario.

BZ#878149

Previously, if no cluster node could have been contacted on certain luci pages, luci displayed the Error 500 message on that page and logged an error message with a traceback into its log. As an appropriate response to this situation, this update modifies luci to display one of the following messages:

```
Unable to contact any of the nodes in this cluster.
```

```
No nodes from this cluster could be contacted. The status of this cluster is unknown
```

BZ#880363

Due to a bug in luci validation code, a confusing validation error message was displayed if a non-existing failover domain in the "Failover Domains" tab was specified. This bug has been fixed and luci now processes these validation errors correctly, displaying appropriate error messages as expected.

BZ#878960

The "User preferences" page was accessible without authentication, which allowed an anonymous user disabling or enabling "expert" mode. Although this behavior had no direct security impact, consistency in assigned authorization is considered to be best practice. This update modifies luci to strictly require users to be authenticated before accessing this "Preferences" page.

BZ#886576

The "Remove this instance" button in the "Edit Fence Instance" form had no function and could have misled cluster administrators. This button has been removed so the aforementioned form now shows only the relevant content.

BZ#1001835

The luci application incorrectly considered the "module_name" parameter of the Dell DRAC 5 fence device as mandatory. Therefore, such a fence device could not have been created without specifying its module name. The validation code has been fixed so luci now treats this parameter as optional, and Dell DRAC 5 fence devices can now be successfully created without module names.

Enhancements**BZ#917814**

A confirmation pop-up dialog has been added that prevents luci from removing selected clusters accidentally.

BZ#983693

The luci application now reflects the concurrent extension to the oracledb, orainstance, and oralistener resource agents regarding Oracle Database 11g support. This also includes the ability to configure the newly supported TNS_ADMIN variable to allow for wider customization.

All luci users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements. After installing this update, the **luci** service will be restarted automatically.

8.108. LVM2**8.108.1. RHBA-2013:1704 — lvm2 bug fix end enhancement update**

Updated lvm2 packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The lvm2 packages include all of the support for handling read and write operations on physical volumes, creating volume groups from one or more physical volumes and creating one or more logical volumes in volume groups.

Bug Fixes**BZ#820991**

When visible clustered volume groups (VGs) were present in the system, it was not possible to silently skip them with proper return error code while the non-clustered locking type was used. To fix this bug, the "--ignoresskippedcluster" option has been added for several LVM commands; namely pvs, vgs, lvs, pvdisplay, vgdisplay, lvdisplay, vgchange, and lvchange. With this option, the clustered VGs are skipped correctly without any warning or error messages while the return error code also does not depend on these clustered VGs.

BZ#834327

Previously, the lvremove command failed to remove a virtual snapshot device if this device was still open. Consequently, the <virtual_snapshot_name>_vorigin device-mapper device was left on the system after the failed removal. A manual remove with use of dmsetup was required to discard this device. With this update, lvremove has been modified to properly check the LV open count status before proceeding with the removal operation.

BZ#861227

Previously, when the `lvconvert` command was used with the `--stripes` option, the required supplementary options, such as `--mirrors` or `--repair`, `thinpool`, or `type raid*/mirror`, were not enforced. Consequently, calling `lvconvert --stripes` without accompanying conversion instructions led to an incomplete conversion. With this update, a condition has been added to enforce the correct syntax. As a result, an error message is now displayed in the described scenario.

BZ#880414

Previously certain `lvm2app` functions were returning values in sectors instead of bytes. This behavior applied for values of `origin_size`, `vg_extent_size`, `stripe_size`, `region_size`, `chunk_size`, `seg_start`, and `pvseg_size`. Consequently, the returned `lvm2app` results were inconsistent and therefore misleading. This behavior has been changed and all `lvm2app` values are now returning byte values.

BZ#902538

The `lvm2` tools determine the PowerPath major number by searching for an `emcpower` line in the `/proc/devices` file. Previously, some versions of PowerPath used the ID string `power2`. As a consequence, on systems with such an identifier, PowerPath devices were not given the expected precedence over PowerPath components which exhibit the same physical volume UUID. With this update, detection of EMC power devices works as expected, and the priority of devices is now set properly.

BZ#902806

Prior to this update, the `lvm2 dmeventd` daemon attempted to reset to C locales only through the `LANG` environmental variable. However, when the system sets locales using the `LC_ALL` variable, this variable has a higher priority than the `LANG` variable, which leads to an extensive memory consumption. With this update, `LC_ALL` has been reset to C instead of `LANG`, thus reducing the memory consumption.

BZ#905254

With this update, a specific diagnostic message has been added for the case when the `lvm2metad` daemon was already running or its pidfile was locked for any other reason. When trying to start `lvm2metad` while it is already running now returns a message with a clear indication of the problem:

```
Failed to acquire lock on /var/run/lvm2metad.pid. Already running?
```

BZ#907487

Previously, the `'vgreduce --removemissing'` command could not be used when missing physical volumes were still used by RAID logical volumes. Now, it is possible for `'vgreduce --removemissing'` to replace the failed physical volume with an `'error'` segment within the affected RAID logical volumes and remove the PV from the volume group. However, in most cases it is better to replace a failed RAID device with a spare one (with use of `'lvconvert --repair'`) if possible.

BZ#910104

Under certain circumstances, cached metadata in the `lvm2metad` daemon could have leaked during metadata updates. With this update, `lvm2metad` has been fixed to prevent the leak.

BZ#913644

Previously, if a device had failed after the `vgexport` command was issued, it was impossible to import the volume group. Additionally, this failure to import also meant it was impossible to repair the volume group. It is now possible to use the `'--force'` option with `vgimport` to import volume groups even if there are devices missing.

BZ#914143

When LVM scans devices for LVM meta data, it applies several filters, such as the multipath filter, MD component filter, or partition signature filter. Previously, the order in which these filters were applied caused that multipath filter failed to filter out a multipath component because the device was accessed by other filters. Consequently, I/O errors occurred if the path was not accessible. With this update, the order of filtering has been changed and the multipath filter now works as expected.

BZ#919604

The 'raid1' type can be used to set the device fault tolerance for thinpool logical volumes. It is no longer possible to create thinpools on top of logical volumes of 'mirror' segment type. The existing thinpools with data or meta data areas of 'mirror' segment type will still function, however, it is recommended to convert these to 'raid1' with use of the 'lvconvert' command.

BZ#928537

When using the pvcreate command with the --restorefile and --uuid options while the supplied UUID was incorrect, an internal error message about a memory leak was issued:

```
Internal error: Unreleased memory pool(s) found.
```

With this update, the memory leak has been fixed and the error message is no longer displayed.

BZ#953612

When updating the device-mapper-event package to a later version, the package update script attempts to restart running dmeventd instance and to replace it with the new dmeventd daemon. However, the previous version of dmeventd does not recognize the notification for restart and therefore a manual intervention is needed in this situation. Previously, the following warning message was displayed:

```
WARNING: The running dmeventd instance is too old
```

In order to provide more precise information and advise for the required action, the following message has been added for the described case:

```
Failed to restart dmeventd daemon. Please, try manual restart
```

BZ#953867

When using the lvmemd daemon together with the accompanying LVM autoactivation feature, the logical volumes on top of encrypted devices were not automatically activated during system boot. This was caused by ignoring the extra udev event that was artificially generated during system boot to initialize all existing devices. This bug has been fixed, and LVM now properly recognizes the udev event used to initialize the devices at boot, including encrypted devices.

BZ#954061

When using the lvmemd daemon together with the accompanying LVM autoactivation feature, the device-mapper devices representing the logical volumes were not refreshed after the underlying PV was unplugged or deactivated and then plugged back or activated. This was caused by assigning a different major and minor pair to identify the reconnected device, while LVs mapped on this device still referenced it with the original pairs. This bug has been fixed and LVM now always refreshes logical volumes on PV device after reactivation.

BZ#962436

Due to a regression introduced in LVM version 2.02.74, when the `optimal_io_size` device hint was smaller than the default `pe_start` size of 1 MiB, this `optimal_io_size` was ignored and the default size was used. With this update, the `optimal_io_size` is applied correctly to calculate the PV's `pe_start` value.

BZ#967247

Prior to this update, before adding additional images to a RAID logical volume, the available space was calculated incorrectly. Consequently, if the available space was insufficient, adding these images failed. This bug has been fixed and the calculation is now performed correctly.

BZ#973519

Previously, if the `nohup` command was used together with LVM commands that do not require input, `nohup` configured the standard input as write-only while LVM tried to reopen it also for reading. Consequently, the commands terminated with the following message:

```
stdin: fdopen failed: Invalid argument
```

LVM has been modified and if the standard input is already open write-only, LVM does not attempt to reopen it for reading.

BZ#976104

Previously, when converting a linear logical volume to a mirror logical volume, the preferred mirror segment type set in the `/etc/lvm/lvm.conf` configuration file was not always accepted. This behavior has been changed, and the segment type specified with the `'mirror_segtype_default'` setting in configuration file is now applied as expected.

BZ#987693

Due to a code regression, a corruption of thin snapshot occurred when the underlying thin-pool was created without the `'--zero'` option. As a consequence, the first 4KB in the snapshot could have been invalidated. This bug has been fixed and the snapshot is no longer corrupted in the aforementioned scenario.

BZ#989347

Due to an error in the LVM allocation code, `lvm2` attempted free space allocation contiguous to an existing striped space. When trying to extend a 3-way striped logical volume using the `lvextend` command, the `lvm2` utility terminated unexpectedly with a segmentation fault. With this update, the behavior of LVM has been modified, and `lvextend` now completes the extension without a segmentation fault.

BZ#995193

Previously, it was impossible to convert a volume group from clustered to non-clustered with a configuration setting of `'locking_type = 0'`. Consequently, problems could arise if the cluster was unavailable and it was necessary to convert the volume group to non-clustered mode. With this update, LVM has been modified to make the aforementioned conversion possible.

BZ#995440

Prior to this update, the repair of inconsistent metadata used an inconsistent code path depending on whether the `lvm2` daemon was running and enabled. Consequently, the `lvm2` version of meta data repair failed to correct the meta data and a warning message was printed repeatedly by every

command until the problem was manually fixed. With this update, the code paths have been reconciled. As a result, metadata inconsistencies are automatically repaired as appropriate, regardless of the `lvmetad`.

BZ#997188

When the `lvm_list_pvs_free` function from the `lvm2app` library was called on a system with no physical volumes, `lvm2app` code tried to free an internal structure that had already been freed before. Consequently, the function terminated with a segmentation fault. This bug has been fixed, and the segmentation fault no longer occurs when calling `lvm_list_pvs_free`.

BZ#1007406

When using LVM logical volumes on MD RAID devices as PVs and while the `lvmetad` daemon was enabled, the accompanying logical volume automatic activation sometimes left incomplete device-mapper devices on the system. Consequently, no further logical volumes could be activated without manual cleanup of the dangling device-mapper devices. This bug has been fixed, and dangling devices are no longer left on the system.

BZ#1009700

Previously, LVM commands could become unresponsive when attempting to read an LVM mirror just after a write failure but before the repair command handled the failure. With this update, a new `'ignore_lvm_mirrors'` configuration option has been added to avoid this issue. Setting this option to `'1'` will cause LVM mirrors to be ignored and prevent the described problem. Ignoring LVM mirrors also means that it is impossible to stack volume groups on LVM mirrors. The aforementioned problem is not present with the LVM RAID types, like `"raid1"`. It is recommended to use the RAID segment types especially when attempting to stack volume groups on top of mirrored logical volumes.

BZ#1016322

Prior to this update, a race condition could occur during the pool destruction in `libdevmapper.so`. Consequently, the `lvmetad` daemon sometimes terminated due to heap corruption, especially under heavier concurrent loads, such as multiple LVM commands executing at once. With this update, a correct locking has been introduced to fix the race condition. As a result, `lvmetad` no longer suffers heap corruption and subsequent crashes.

BZ#1020304

The `blkdeactivate` script iterates over the list of devices given to it as an argument and tries to unmount or deactivate them one by one. However, in case of failed unmount or deactivation, the iteration did not proceed. Consequently, `blkdeactivate` kept attempting to process the same device and entered an endless loop. This behavior has been fixed and if `blkdeactivate` fails to unmount or deactivate any of the devices, the processing of this device is properly skipped and `blkdeactivate` proceeds as expected.

Enhancements**BZ#814737**

With this update, `lvm2` has been enhanced to support the creation of thin snapshots of existing non-thinly-provisioned logical volumes. Thin-pool can now be used for these snapshots of non-thin volumes, providing performance gains. Note that the current `lvm2` version does not support the merge feature, so unlike with older `lvm2` snapshots, an updated device cannot be merged back into its origin device.

BZ#820203

LVM now supports validating of configuration files and it can report any unrecognized entries or entries with wrong value types in addition to existing syntax checking. To support this feature, a new "config" configuration section has been added to the /etc/lvm/lvm.conf configuration file. This section has two configurables: "config/checks" which enables or disables the checking (enabled by default), and "config/abort_on_errors" which enables or disables immediate abort on any invalid configuration entry found (disabled by default).

In addition, new options have been added to the "lvm dumpconfig" command that make use of the new configuration handling code introduced. The "lvm dumpconfig" now recognizes the following options: --type, --atversion, --ignoreadvanced, --ignoreunsupported, --mergedconfig, --withcomments, --withversions, and --validate.

BZ#888641

Previously, the scm (Storage Class Memory) device was not internally recognized as partitionable device. Consequently, scm devices could not be used as physical volumes. With this update, scm device has been added to internal list of devices which are known to be partitionable. As a result, physical volumes are supported on scm partitions. Also, the new 'lvm devtypes' command has been added to list all known device types.

BZ#894136

When the lvm2 daemon is enabled, meta data is cached in RAM and most LVM commands do not consult on-disk meta data during normal operation. However, when meta data becomes corrupt on disk, LVM may not take a notice until a restart of lvm2 or a reboot. With this update, the vgck command used for checking VG consistency has been improved to detect such on-disk corruption even while lvm2 is active and the meta data is cached. As a result, users can issue the "vgck" command to verify consistency of on-disk meta data at any time, or they can arrange a periodic check using cron.

BZ#903249

If a device temporarily fails, the kernel notices the interruption and regards the device as disabled. Later, the kernel needs to be notified before it accepts the device as alive again. Previously, LVM did not recognize these changes and the 'lvs' command reported the device as operating normally even though the kernel still regarded the device as failed. With this update, 'lvs' has been modified to print a 'p' (partial) if a device is missing and also an 'r' (refresh/replace) if the device is present but the kernel regards the device as still disabled. When seeing an 'r' attribute for a RAID logical volume, the user can then decide if the array should be refreshed (reloaded into the kernel using 'lvchange --refresh') or if the device should be replaced.

BZ#916746

With this update, snapshot management handling of COW device size has been improved. This version trims the snapshot COW size to the maximal usable size to avoid unnecessary disk space consumption. It also stops snapshot monitoring once the maximal size is reached.

BZ#921280

Support for more complicated device stack for thinpool has been enhanced to properly resize more complex volumes like mirrors or raids. The new lvm2 version now supports thin data volume extension on raids. Support for mirrors has been deactivated.

BZ#921734

Prior to this update, the "vgchange -c {y|n}" command call changed all volume groups accessible on the system to clustered or non-clustered. This may have caused an unintentional change and therefore the following prompt has been added to acknowledge this change:

Change clustered property of all volumes groups? [y/n]

This prompt is displayed only if the "vgchange -c {y|n}" is called without specifying target volume groups.

BZ#924137

The blkdeactivate utility now suppresses error and information messages from external tools that are called. Instead, only a summary message "done" or "skipped" is issued by blkdeactivate. To show these error messages if needed, a new -e/--errors switch has been added to blkdeactivate. Also, there's a new -v/--verbose switch to display any information messages from external tools together with any possible debug information.

BZ#958511

With this update, the blkdeactivate utility has been modified to correctly handle file systems mounted with bind (the 'mount -o bind' command). Now, blkdeactivate unmounts all such mount points correctly before trying to deactivate the volumes underneath.

BZ#969171

When creating many RAID logical volumes at the same time, it is possible for the background synchronization I/O necessary to calculate parity or copy mirror images to crowd out nominal I/O and cause subsequent logical volume creation to slow dramatically. It is now possible to throttle this initializing I/O via the '--raidmaxrecoveryrate' option to lvcreate. You can use the same argument with lvchange to alter the recovery I/O rate after a logical volume has been created. Reducing the recovery rate will prevent nominal I/O from being crowded out. Initialization will take longer, but the creation of many logical volumes will proceed more quickly. (BZ#969171)

BZ#985976

With this update, RAID logical volumes that are created with LVM can now be checked with use of scrubbing operations. Scrubbing operations are user-initiated checks to ensure that the RAID volume is consistent. There are two scrubbing operations that can be performed by appending the "check" or "repair" option to the "lvchange --syncaction" command. The "check" operation will examine the logical volume for any discrepancies, but will not correct them. The "repair" operation will correct any discrepancies found.

BZ#1003461

This update adds support for thin external origin to lvm2. This allows to use any LV as an external origin for a thin volume. All unprovisioned blocks are loaded from the external origin volume, while all once-written blocks are loaded from the thin volume. This functionality is provided by the 'lvcreate --snapshot' command and the 'lvconvert' command that converts any LV into a thin LV.

BZ#1003470

The error message 'Cannot change discards state for active pool volume "pool volume name"' has been improved to be more comprehensible: 'Cannot change support for discards while pool volume "pool volume name" is active'.

BZ#1007074

The repair of corrupted thin pool meta data is now provided by the 'lvconvert --repair' command, which is low-level manual repair. The thin pool meta data volume can be swapped out of the thin-pool LV via 'lvconvert --poolmetadata swapLV vg/pool' command and then the thin_check, thin_dump, and

thin_repair commands can be used to run manual recover operation. After the repair, the thin pool meta data volume can be swapped back. This low-level repair should be only used when the user is fully aware of thin-pool functionality.

BZ#1017291

LVM now recognizes NVM Express devices as a proper block device type.

Users of lvm2 are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.109. MAILX

8.109.1. RHBA-2013:1129 — mailx bug fix update

Updated mailx packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The mailx packages contain a mail user agent that is mostly used to manage mail using scripts.

Bug Fixes

BZ#845098

Prior to this update, the "mail" utility provided with the Red Hat Enterprise Linux 6 mailx packages was not fully compatible with the utility provided with the Red Hat Enterprise Linux 5 mailx package and packages in earlier releases of Red Hat Enterprise Linux. Consequently, some user scripts written for the mail utility did not work with "mail" in Red Hat Enterprise Linux 6. Support for multiple versions of the "mail" utility has been added to the mailx packages. This allows the user to install alternative packages providing this utility, for example, bsd-mailx.

BZ#857120

The mailx command did not set the error return code when it failed to send an e-mail because the TMPDIR environment variable was set to an invalid path. As a consequence, error checking was incorrect and therefore not helpful. With this update, the correct return code is set when mailx fails to send an e-mail. The error checking now works properly.

Users of mailx are advised to upgrade to these updated packages, which fix these bugs.

8.110. MAN-PAGES-FR

8.110.1. RHBA-2013:1093 — man-pages-fr bug fix update

Updated man-pages-fr packages that fix one bug are now available.

The man-pages-fr packages contain manual pages in French.

Bug Fix

BZ#903048

Due to some problem in the build system of the French manual page package man-pages-fr, some manual pages were not included in the package. Some manual pages, for example the manual page of "echo" were displayed in English even when the system was running in a French locale. Thus, the

command “man echo” displayed an English manual page. The build problem in the man-pages-fr package is fixed, and the missing manual pages are now included. Hence, manual pages are now displayed in French when the system is running in a French locale, for example “man echo” now shows a French manual page.

Users of man-pages-fr are advised to upgrade to these updated packages, which fixes this bug.

8.111. MAN-PAGES-JA

8.111.1. [RHBA-2013:1094 — man-pages-ja bug fix update](#)

An updated man-pages-ja package that fixes two bugs is now available for Red Hat Enterprise Linux 6.

The man-pages-ja package contains manual pages in Japanese.

Bug Fixes

BZ#949787

The shmat(2) man page in the previous release did not mention the EIDRM error code, which could have been returned by the shmat utility. With this update, the EIDRM error code is included in shmat.

BZ#957937

The strtoul(3) man page in the previous release incorrectly mentioned the range of the return value. This update fixes the aforementioned problem.

Users of man-pages-ja are advised to upgrade to this updated package, which fixes these bugs.

8.112. MAN-PAGES-OVERRIDES

8.112.1. [RHBA-2013:1695 — man-pages-overrides bug fix and enhancement update](#)

Updated man-pages-overrides packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The man-pages-overrides packages provide a collection of manual (man) pages to complement other packages or update those contained therein.

Bug Fixes

BZ#988125

The madvise(2) manual page did not contain the "MADV_DODUMP" and "MADV_DONTDUMP" arguments. This update adds a description of these arguments to the madvise(2) manual page.

BZ#833868

Previously, the manual page for the dig utility contained upstream-specific options for an Internationalized Domain Name (IDN) library. Consequently, these options did not function as expected and users were incapable of disabling IDN support in dig following the steps from the

manual page. The dig(1) manual page has been modified to include the options of the IDN library used in Red Hat Enterprise Linux and users can now successfully disable IDN support in dig following the steps from the manual page.

BZ#978981

Previously, no manual page for the getent utility was available in Red Hat Enterprise Linux 6. This update adds the missing getent(1) manual page and the documentation of the utility is now complete.

BZ#872144

Prior to this update, the top(1) manual page did not describe calculation of resident memory size properly. The incorrect calculation of resident memory size has been removed from this manual page.

BZ#1018622

Previously, the description of the "new station" message in the arpwatch(8) manual page was not accurate, which could cause confusion. This update adds a correct description of the "new station" message to the arpwatch(8) manual page.

BZ#896700

Previously, the auditd.conf(5) manual page contained an incomplete sentence. The incomplete sentence has been fixed with this update.

BZ#974697

The ld.so(8) manual page contained an incorrect description of the "LD_PRELOAD" variable. With this update, the description of the variable has been corrected in the ld.so(8) manual page.

BZ#1002071

The bash(1) manual page did not reflect the behavior changes in the "extglob" option introduced in Bash version 4.1. This update adds a correct description of "extglob" behavior to the bash(1) manual page.

BZ#903258

The manual page for the fallocate utility did not contain description of the "FALLOC_FL_PUNCH_HOLE" flag. This update adds a description of "FALLOC_FL_PUNCH_HOLE" to the fallocate(2) manual page.

BZ#979318

Previously, the manual page for the netstat utility did not mention IPv6 in the description of the command's "-A" option. With this update, the description of the IPv6 functionality has been added to the netstat(8) manual page.

BZ#905066

Previously, loading Certification Authority (CA) certificates by nickname when using the curl utility with Network Security Services (NSS) was described incorrectly in the curl documentation. This update adds correct documentation of the above mentioned process.

BZ#957010

Prior to this update, the strtoul(3) manual page contained incorrect return values of the "strtoul()" and "strtoull()" functions. This update fixes the strtoul(3) manual page and it now contains correct information.

BZ#960281

Previously, the `clock_getres(2)` manual page did not contain the "CLOCK_MONOTONIC_COARSE" and "CLOCK_REALTIME_COARSE" `clk_id` values. With this update, the above mentioned values have been added to the `clock_getres(2)` manual page.

BZ#974685

Previously, the `sched_setaffinity(2)` manual page contained an incorrect example, which could cause confusion. The incorrect example has been removed from the `sched_setaffinity(2)` manual page.

BZ#951826

Previously, the manual page for the `postconf` utility contained incorrect information about the default configuration of a postfix server. This update fixes the default configuration description in the `postconf(5)` manual page.

BZ#979460

The `mailx(1)` manual page contained an incomplete entry about setting variables with the "-S" option. This update provides a better description of setting variables with the "-S" option and adds an example of the syntax using "from=" to the `mailx(1)` manual page.

BZ#913191

Prior to this update, the `selinux(8)` manual page contained outdated information. This manual page has been updated, and SELinux is now documented correctly.

BZ#907837

The `useradd(8)` manual page contained an incorrect description of the "-u, --uid UID" option, which could cause confusion. The description has been fixed with this update.

BZ#807323

Previously, a manual page for the `byzanz-record` utility did not mention the possibility to use the "webm" output format and the manual page was thus incomplete. This update adds "webm" to the `byzanz-record(1)` manual page.

BZ#1020417

The `ssh_config(5)` manual page contained an incorrect default value for the "KexAlgorithms" option. This bug has been fixed and the default value for "KexAlgorithms" in the `ssh_config(5)` manual page is now correct.

BZ#1020432

The "-n" option in the `ssh-keygen` utility was renamed to "-Z", but the `ssh-keygen(1)` manual page was not updated. This bug has been fixed and the `ssh-keygen(1)` manual page now describes the correct option.

Enhancement**BZ#928917**

Previously, the "O_DIRECT" flag was not described in the `open(2)` manual page. This update adds this description and the documentation is now complete.

Users of man-pages-overrides are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.113. MCELOG

8.113.1. RHBA-2013:1658 — mcelog bug fix and enhancement update

Updated mcelog packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The mcelog packages contain a daemon that collects and decodes Machine Check Exception (MCE) data on AMD64 and Intel 64 machines.

Bug Fixes

BZ#875824

Previously, mcelog packages installed a cron job to report the status of mce logs, which conflicted with running the mclogd service as default mode. Consequently, mcelog competed with the cron job and did not collect complete data. With this update, cron job is not installed in case mcelogd is running, thus fixing this bug.

BZ#919999

Due to a bug in mcelog packages, the AMD Family 15 architecture was not supported. The bug has been fixed and mcelog now supports AMD Family 15 as expected.

BZ#996634

Previously, support for extended logging was enabled by default in mcelog packages. Consequently, on systems with processors without support for extended logging, the mcelog service terminated unexpectedly with the following message:

```
mcelog: Cannot open /dev/cpu/0/msr to set imc_log: Permission denied
```

With this update, extended logging is disabled by default in mcelog packages, and the mcelog service no longer crashes in the aforementioned scenario.

Enhancement

BZ#881555, BZ#922873, BZ#991079

With this update, mcelog packages support Intel Xeon Processor E5-XXXX v3, Intel Xeon Processor E5-XXXX, and Intel Xeon Processor E3-XXXX v3 architectures.

Users of mcelog are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.114. MDADM

8.114.1. RHBA-2013:1643 — mdadm bug fix and enhancement update

Updated mdadm packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The mdadm packages contain a utility for creating, managing, and monitoring Linux multiple disk (MD) devices.



NOTE

The mdadm packages have been upgraded to upstream version 3.2.6, which provides a number of bug fixes and enhancements over the previous version, including performance improvements. (BZ#[922971](#))

Bug Fixes

BZ#[903212](#)

Previously, during expanding the size of an Intel Matrix Storage Manager (IMSM) RAID1 or RAID5 volume, the resynchronization process was reported in the `/proc/mdstat` file but there was no information about the process stored in the volume's metadata. Consequently, if the RAID volume was stopped during the process of size expansion, all information about this progress was lost and the resynchronization would be restarted from the beginning on the next array reassembly. A patch has been applied to address this problem, and information is now stored in metadata as expected in the described scenario.

BZ#[950545](#)

Prior to this update, the mdadm utility did not work correctly when attempting to write a superblock onto a defective drive. Consequently, mdadm could terminate unexpectedly with a segmentation fault if it encountered a write error. This bug has been fixed and mdadm no longer crashes in this scenario.

BZ#[955972](#)

Previously, the mdadm utility did not work correctly if a rebuild of an Intel Matrix Storage Manager (IMSM) RAID5 volume was started in Option ROM (OROM). Consequently, the RAID5 volume was in the "degraded" state once booted into the operating system and the rebuild did not proceed. A patch has been applied to address this problem and rebuilding IMSM RAID5 volumes now completes successfully in the described scenario.

BZ#[956016](#)

Previously, when an Intel Matrix Storage Manager (IMSM) volume was being reshaped, the "mdadm -Ss" command used for stopping the process did not work properly. Consequently, on the first run of "mdadm -Ss", only the volume was stopped but the container was left in place, and a second execution of the command was necessary. This bug has been fixed and the command now works as expected during a volume's reshape.

BZ#[995105](#)

Previously, when an Intel Matrix Storage Manager (IMSM) RAID10 volume was being resynchronized or rebuilt, stopping the process after 50% completion did not work properly. As a consequence, the processes did not proceed correctly after reassembling, and the data became corrupted. With this update, resynchronization and rebuild work correctly in this scenario.

BZ#[1001627](#)

Prior to this update, when an Intel Matrix Storage Manager (IMSM) RAID1 or RAID10 volume was being rebuilt and this process was stopped, an attempt to resume the rebuild was not successful. Consequently, the rebuild did not start even when a new drive was added to the container, and metadata contained incorrect information. This bug has been fixed and resuming a rebuild now works properly in the described scenario.

BZ#1010859

Previously, the mdadm utility did not work correctly when a disk failed in an Intel Matrix Storage Manager (IMSM) RAID volume. Consequently, the failed disk was removed neither from the volume nor from the container, the volume was not in the "degraded" state, and the rebuild could not start. With this update, mdadm handles failed disks in RAID volumes properly.

Users of mdadm are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.115. MESA

8.115.1. RHBA-2013:1559 — mesa bug fix and enhancement update

Updated mesa packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

Mesa provides a 3D graphics API that is compatible with Open Graphics Library (OpenGL). It also provides hardware-accelerated drivers for many popular graphics chips.

Bug Fixes

BZ#879637

On certain Intel GT2+ processors, segmentation faults could have been reported in the output of the dmesg command after running a Piglit quick-driver test. A patch has been applied to address his bug, and the unwanted behavior no longer occurs.

BZ#908547

Prior to this update, compressed texture size checks were performed in an incorrect manner. Consequently, checking the image size against the compression block size could cause certain applications to terminate unexpectedly. The underlying source code has been modified, and the texture error no longer causes the applications to crash in the described scenario.

Enhancements

BZ#818345

Support for future Intel 2D and 3D graphics has been added to allow systems using future Intel processors to be certified through the Red Hat Hardware Certification program.

BZ#957792

With this update, the mesa-private-llvm library has been added.

Users of mesa are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.116. MICROCODE_CTL

8.116.1. RHBA-2013:1668 — microcode_ctl bug fix and enhancement update

Updated `microcode_ctl` packages that fix one bug and add one enhancement are now available for Red Hat Enterprise Linux 6.

The `microcode_ctl` packages provide utility code and microcode data to assist the kernel in updating the CPU microcode at system boot time. This microcode supports all current x86-based, Intel 64-based, and AMD64-based CPU models. It takes advantage of the mechanism built-in to Linux that allows microcode to be updated after system boot. When loaded, the updated microcode corrects the behavior of various processors, as described in processor specification updates issued by Intel and AMD for those processors.

Bug Fix

BZ#1000317

Previously, the `microcode_ctl` utility did not detect if it was running in a virtual machine and attempted to install the CPU microcode updates. This behavior caused several errors to be returned in the kernel ring buffer. The underlying source code has been modified and `microcode_ctl` no longer tries to update the CPU microcode in the described scenario.

Enhancement

BZ#915957, BZ#1005606

The Intel CPU microcode file has been updated to version 20130906.

All users of `microcode_ctl` are advised to upgrade to these updated packages, which fix this bug and add this enhancement. Note: a system reboot is necessary for this update to take effect.

8.117. MOBILE-BROADBAND-PROVIDER-INFO

8.117.1. RHBA-2013:0974 — mobile-broadband-provider-info bug fix update

Updated `mobile-broadband-provider-info` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `mobile-broadband-provider-info` packages contain listings of mobile broadband (3G) providers, associated network, and plan information.

Bug Fix

BZ#844288

Previously, in the `serviceproviders.xml` file located in the `/usr/share/mobile-broadband-provider-info/` directory, "internet.saunalahti" was incorrectly specified as an APN (Access Point Name) value for the Sonera provider. This prevented the Sonera mobile broadband configuration from working. The stanza containing "internet.saunalahti" as an APN value for Sonera has been removed from the XML file, and the Sonera mobile broadband configuration now works as expected.

Users of `mobile-broadband-provider-info` are advised to upgrade to these updated packages, which fix this bug.

8.118. MOD_AUTH_KERB

8.118.1. RHBA-2013:0860 — [mod_auth_kerb bug fix update](#)

Updated `mod_auth_kerb` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `mod_auth_kerb` package provides a module for the Apache HTTP Server designed to provide Kerberos authentication over HTTP. The module supports the Negotiate authentication method, which performs full Kerberos authentication based on ticket exchanges.

Bug Fix

BZ#867153

Previously, when the `KrbLocalUserMapping` directive was enabled, `mod_auth_kerb` did not translate a principal name properly if the local name was of a higher length. Consequently, the Apache server returned the HTTP 500 error in such a scenario. A patch has been provided to address this issue and the module now correctly translates account names longer than their counterpart principal names.

Users of `mod_auth_kerb` are advised to upgrade to these updated packages, which fix this bug.

8.119. MODEMMANAGER

8.119.1. RHBA-2013:1672 — [ModemManager bug fix update](#)

Updated `ModemManager` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `ModemManager` packages provide a consistent application programming interface (API) to operate a wide variety of modems, including mobile broadband (3G) devices.

Bug Fix

BZ#883079

Previously, some broadband devices were not covered by the "udev" rules in the `/lib/udev/rules.d/77-mm-*.rules` files. As a consequence, the broadband connection was either not established at all, or failed after communicating a few packages. Additional "udev" rules have been included to ensure that `ModemManager` uses the correct serial port. As a result, more broadband devices, such as ZTE, LG, and Sierra Wireless modems, are now supported.

Users of `ModemManager` are advised to upgrade to these updated packages, which fix this bug.

8.120. MYSQL

8.120.1. RHBA-2013:1647 — [mysql bug fix update](#)

Updated `mysql` packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

MySQL is a multi-user, multi-threaded SQL database server. It consists of the MySQL server daemon (`mysqld`) and many client programs and libraries.

Bug Fixes

BZ#842052

Prior to this update, the `mysqld` daemon worked with uninitialized memory when accessing non-

nullable GEOMETRY types. Consequently, mysqld could terminate unexpectedly when the mysqldump utility was running. With this update, mysqld initializes memory properly and thus no longer crashes in this scenario

BZ#877557

Previously, the mysqldump utility expected log tables to be created on the MySQL 5.0.x server, from which it retrieved data. Consequently, mysqldump could not dump the MySQL system table. With this update, mysqldump no longer expects log tables to be created, and it is now able to dump the system table in the described scenario as expected.

BZ#884651

Prior to this update, the mysqld init script did not correctly verify the status of the mysqld daemon. Consequently, the script could return an error message even when the daemon had successfully started. The mysqld init script has been fixed, and it now checks the daemon status properly.

BZ#904061

Previously, the mysql-server sub-packages did not contain the logrotate script. Consequently, the log rotation had to be configured manually. With this update, the logrotate script has been provided by the mysql-server sub-packages, and users can use the script to log into the mysqld.log file by uncommenting appropriate lines in the script.

Users of mysql are advised to upgrade to these updated packages, which fix these bugs. After installing this update, the MySQL server daemon (mysqld) will be restarted automatically.

8.121. NET-SNMP

8.121.1. RHBA-2013:1693 — net-snmp bug fix and enhancement update

Updated net-snmp packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The net-snmp packages provide a generic client library, a suite of command-line tools, an extensible SNMP agent, Perl modules, and Python modules to use and deploy the Simple Network Management Protocol (SNMP).

Bug Fixes

BZ#893119

Previously, snmpd, the SNMP daemon, did not check for errors when populating data for the UCD-SNMP-MIB::extTable table and could leak memory when the system ran out of memory. This bug has been fixed and snmpd now checks for out-of-memory conditions and frees the memory for the UCD-SNMP-MIB::extTable table when encounters an error.

BZ#907571

Previously, the snmp_config(5) manual page was not clear about which files were looked for and the reader could get the incorrect impression that any file with a suffix "conf" or "local.conf" could be used as an snmp configuration file. In this update, the snmp_config(5) manual page has been modified to precisely specify which files are used as snmp configuration files.

BZ#919259

In a previous update, the `snmpd` daemon was fixed to show the executable name and all the command-line arguments in the `UCD-SNMP-MIB::extCommand` OID string. The fix did not check for executables without command-line arguments. Consequently, the `snmpd` daemon terminated unexpectedly with a segmentation fault when retrieving the value of the `UCD-SNMP-MIB::extCommand` OID of an executable with no arguments. With this update, `snmpd` now checks if there are no arguments and shows the correct value of the `UCD-SNMP-MIB::extCommand` OID. As a result, crashes no longer occur in the described scenario.

BZ#919952

In previous `net-snmp` package updates, the `HOST-RESOURCES-MIB::hrSWRunTable` table was rewritten, and, due to a regression, it did not report the `"hrSWRunPath"` string of kernel threads. This update fixes the `HOST-RESOURCES-MIB::hrSWRunPath` string of kernel threads and is now reported by the `snmpd` daemon.

BZ#922691

When the `"includeAllDisks"` configuration option was specified in the `/etc/snmp/snmpd.conf` file, the `snmpd` daemon scanned the running system only at startup and did not update the `UCD-SNMP-MIB::dskTable` table if a new device was mounted later. As a consequence, on dynamic systems where devices are frequently mounted and unmounted, `UCD-SNMP-MIB::dskTable` could not be used to monitor storage usage, because it monitored only devices which were available at system start. To fix this bug, the implementation of `UCD-SNMP-MIB::dskTable` was enhanced to dynamically add new devices as they are mounted. This happens only when the `"includeAddDisks"` configuration option is used in `/etc/snmp/snmpd.conf`. As a result, in dynamic systems where devices are frequently mounted and unmounted, `UCD-SNMP-MIB::dskTable` always shows the current list of mounted devices.

BZ#927474

Previously, `snmpd`, the SNMP daemon, did not set a proper message size when communicating with the Linux kernel using a netlink socket. As a consequence, the message `"netlink: 12 bytes leftover after parsing attributes."` was saved to the kernel log. With this update, `snmpd` sets a correct message size and the kernel no longer logs the aforementioned message.

BZ#947973

In previous Net-SNMP releases, `snmpd` reported an invalid speed of network interfaces in `IF-MIB::ifTable` and `IF-MIB::ifXTable` tables if the interface had a speed other than 10, 100, 1000 or 2500 MB/s. Thus, the returned `net-snmp ifHighSpeed` value was `"0"` compared to the correct speed as reported in `ethtool`, if the Virtual Connect speed was set to, for example, 0.9 Gb/s. With this update, the `ifHighSpeed` value returns the correct speed as reported in the `ethtool` utility, and `snmpd` correctly reports non-standard network interface speeds.

BZ#953926

Net-SNMP did not verify if incoming SNMP messages were encoded properly. In some instances, it read past the receiving buffer size when parsing a message with an invalid size of an integer filed in the message. This caused `snmptrapd`, the SNMP trap processing daemon, to terminate unexpectedly with a segmentation fault on the incoming malformed message. This update enhances the checks of incoming messages and `snmptrapd` no longer crashes when parsing incoming messages with invalid integer sizes.

BZ#955771

Previously, the Net-SNMP python module did not propagate various errors to applications which use this module. As a consequence, the applications were not aware of errors, which had occurred during the SNMP communication. To fix this bug, the Net-SNMP python module has been updated to return

the proper error codes. As a result, the applications now receive information about SNMP errors.

BZ#960568

In previous releases, the `snmp-bridge-mib` subagent included the bridge itself as a port of the bridge in the `BRIDGE-MIB::dot1dBasePortTable` table. This bug has been fixed and the `snmp-bridge-mib` subagent now reports only real interfaces as ports in the `BRIDGE-MIB::dot1dBasePortTable` table.

BZ#968898

Previously, the `snmpd` daemon did not properly terminate strings when processing the "agentaddress" configuration option. As a consequence, when the configuration was re-read multiple times using the `SIGHUP` signal, a buffer overflow occurred. This bug has been fixed and `snmpd` now properly terminates strings during an "agentaddress" processing and no longer crashes using the `SIGHUP` signal.

BZ#983116

The previous Net-SNMP update contained a fix to improve the checking of invalid incoming SNMP messages. This fix introduced a regression and some valid SNMP messages with multiple variables inside were marked as invalid. As a consequence, Net-SNMP tools and servers rejected valid SNMP messages and waited for a "proper" response until timeout. With this update, valid SNMP messages are no longer rejected. As a result, the servers and utilities accept the first incoming message and do not wait for a timeout.

BZ#989498, BZ#1006706

In the previous Net-SNMP updates, the implementation of the `HOST-RESOURCES-MIB::hrStorageTable` table was rewritten and devices with Virtuozzo File System (VZFS) and B-tree File System (BTRFS) were not reported. After this update, `snmpd` properly recognizes devices using VZFS and BTRFS file systems and reports them in `HOST-RESOURCES-MIB::hrStorageTable`.

BZ#991213

Previously the `snmpd` daemon incorrectly parsed Sendmail configuration files with enabled queue groups. Consequently, `snmpd` entered a loop on startup. This update fixes the parsing of configuration files with queue groups and `snmpd` no longer enters a loop on startup.

BZ#1001830

Previously, the Net-SNMP utilities and daemons blindly expected that an MD5 hash algorithm and a DES encryption were available in the system's OpenSSL libraries and did not check for errors when using these cryptographic functions. As a consequence, the Net-SNMP utilities and daemons terminated unexpectedly when attempting to use an MD5 or DES algorithm which are not available when the system is running in FIPS mode. The Net-SNMP utilities and daemons now check for cryptographic function error codes and display the following error message:

```
Error: could not generate the authentication key from the supplied pass
phrase
```

As a result, the aforementioned utilities and daemons no longer crash in FIPS mode.

Enhancements

BZ#917816

After this update, all net-snmp configuration files can use the "includeFile" and "includeDir" options to include other configuration files or whole directories of configuration files. Detailed syntax and usage is described in the snmp_config(5) manual page.

BZ#919239

Previously, the Net-SNMP application was shipping its configuration files, which could contain sensitive information like passwords, readable to any user on the system. After this update, the configuration files are readable only by the root user.

Users of net-snmp are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.122. NETCF

8.122.1. RHBA-2013:1660 — netcf bug fix update

Updated netcf packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The netcf packages contain a library for modifying the network configuration of a system. Network configuration is expressed in a platform-independent XML format, which netcf translates into changes to the system's "native" network configuration files.

Bug Fixes

BZ#844578

When using the "virsh iface-start" or "nctool ifup" command to start a disconnected interface configured to use a DHCP server, the netcf library reported a failure. However, the subsequent list of all interfaces showed the interface as "active". After this update, netcf only reports "active" interface status when the interface is marked both "UP" and "RUNNING" by the "ifconfig" utility and if any attempt to start the interface was successful.

BZ#848722

Previously, attempts to define an interface with a netmask higher than 24 bits failed. This bug has been fixed and it is now possible to define interfaces with netmasks of up to 30 bits.

Users of netcf are advised to upgrade to these updated packages, which fix these bugs.

8.123. NETWORKMANAGER

8.123.1. RHBA-2013:1670 — NetworkManager bug fix and enhancement update

Updated NetworkManager packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

NetworkManager is a system network service that manages network devices and connections, attempting to keep network connectivity active when available. It manages Ethernet, Wi-Fi, mobile broadband (WWAN), and PPPoE (Point-to-Point Protocol over Ethernet) devices, and provides integration with a variety of VPN services.

Bug Fixes

BZ#922558

Previously, **NetworkManager** did not explicitly request static routes from DHCP (Dynamic Host Configuration Protocol) servers, and thus some servers would not deliver those routes. With this update, **NetworkManager** now requests static routes from DHCP servers when available.

BZ#701381

Previously, it was impossible for some users to check **Enable Wireless** box in **NetworkManager** as the field was unresponsive. Moreover, the **Enable Wireless** connection option was unavailable in **NetworkManager** after hardware was disabled and enabled again. With this update, users can turn on the wireless connection from the GUI after their hardware is reenabled.

BZ#1008884

When running the **NetworkManager** applet in some Virtual Machine (VM) configurations, left-clicking on the icon could cause the applet to terminate unexpectedly. This bug has been fixed and the applet no longer crashes in these configurations.

BZ#923648

Previously, bridge and bond connections created through the **NetworkManager** connection editor (**nm-connection-editor**) were not set to connect automatically, and thus had to be manually started. With this update, these connections automatically start when created by default.

BZ#896198

A **GATEWAY** setting in the `/etc/sysconfig/network` file caused **NetworkManager** to assign that **GATEWAY** to all interfaces with static IP addresses. This scenario took place even if no **GATEWAY** or a different one was specified for these addresses. To fix this bug, if **GATEWAY** is given in `/etc/sysconfig/network`, only configurations with a matching gateway address will be given the default route. Alternatively, the **DEFROUTE=yes/no** option may be used in individual configuration files to allow or deny the default route on a per-configuration basis.

BZ#836993

Previously, when using the **vpnc** program via **NetworkManager** with token out of synchronization, the server prompted for a next token. However, **NetworkManager** misinterpreted this response and reported a failed connection. With this update, a new prompt for next token code has been added to the **NetworkManager -vpnc** utility, thus fixing the bug.

BZ#991341

Prior to this update, on receipt of an IPv6 Router Advertisement, **NetworkManager** attempted to replace the IPv6 default route which the kernel had added. Consequently, the kernel returned the following failure message:

```
'ICMPv6 RA: ndisc_router_discovery() failed to add default route.'
```

To fix this bug, **NetworkManager** no longer replaces an IPv6 default route added by the kernel.

BZ#758076

Previously, it was not possible to choose Certificate Authority (CA) certificate via the "Choose certificate" dialog window in **nm-connection-editor**. This was confusing for the user. The dialog checkbox information has been replaced with a more informative text, thus fixing the bug.

BZ#919242

Previously, when **NetworkManager** was not allowed to manage bridge, bond, or VLAN interfaces due to the missing `NM_BOND_BRIDGE_VLAN_ENABLED` option in the `/etc/sysconfig/network` file, the **NetworkManager** connection editor (`nm-connection-editor`) still allowed the user to create these types of network connections. The editor now warns the user when unusable connections have been created, thus fixing the bug.

BZ#915480

Previously, the **NetworkManager** GUI applet (`nm-applet`) did not show bridge, bond, or VLAN interfaces in the menu. With this update, the `nm-applet` has been enhanced to show all available bond, bridge, and VLAN interfaces that are configured but not yet created.

BZ#905532

Due to some missing ignored options for bonding interfaces, the `/sys/class/net/bond0/bonding/primary` file was empty during installation. In addition, the network traffic went through `eth0` during installation. This bug has been fixed and **NetworkManager** now supports a much larger set of bond interface options.

BZ#953076

Previously, in some cases, **NetworkManager** was unable to set the mode of a bond master interface. A patch has been provided to fix this bug and the mode setting now changes according to `nm-editor` alterations.

BZ#953123

Previously, the **NetworkManager** connection editor (`nm-connection-editor`) did not allow setting the cloned MAC address for VLAN interfaces. A patch has been provided to fix this bug and `nm-connection-editor` now works as expected.

BZ#969363

Prior to this update, the manual page of `nm-online` did not describe the correct usage of `nm-online` parameters, such as the `-t` option. The manual page has been updated to describe the usage of its parameters correctly.

BZ#973245

Previously, **NetworkManager** wrote and saved only connection types compatible with standard `ifcfg` network configuration files. This bug has been fixed and other connection types like Bluetooth, WWAN, can now be saved as keyfiles in the `/etc/NetworkManager/system-connections/` directory.

BZ#902372

Previously, when taking control of an existing bridge, **NetworkManager** did not ensure a clean bridge state. With this update, **NetworkManager** resets bridge options and removes all bridge ports, which ensures clean bridge state on start-up with bridging support enabled.

BZ#867273

After configuring the IP-over-InfiniBand (IPoIB) profile on machine with an InfiniBand (IB) device, the profile was not connected. This bug has been fixed and IP-over-Infiniband (IPoIB) network configurations are now listed in the network applet menu.

BZ#713975

After changing the authentication or inner authentication drop-down menus in the configuration for a new wireless network connection, the "Ask for this password every time" checkbox kept resetting. To fix this bug, the updated **NetworkManager** GUI applet saves the value of the checkbox when connecting to WPA Enterprise networks.

BZ#906133

Prior to this update, an Ad-Hoc WiFi network failed to start when its BSSID (Basic Service Set Identifier) was specified, due to kernel restrictions. To fix this bug, the **NetworkManager** connection editor (**nm-connection-editor**) disallows setting the BSSID for ad-Hoc WiFi connections, since this value is automatically chosen by the kernel.

Enhancements**BZ#602265**

With this update, **NetworkManager** has been enhanced to support the creation and management of Point-to-point Protocol over Ethernet (PPPoE) based connections. **NetworkManager** now waits a short period of time before reconnecting a PPPoE connection to ensure the peer is ready.

BZ#694789

A new **GATEWAY_PING_TIMEOUT** configuration option has been added. This new option ensures that **NetworkManager** waits for a successful ping of the **gateway** before indicating network connectivity.

BZ#990310

NetworkManager now reads **ifcfg** alias files and assigns the addresses in them to their master interface, using the alias name as the address label.

BZ#564467, BZ#564465

Manual pages for **nm-connection-editor** and **nm-applet** utilities have been created.

Users of NetworkManager are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.124. NFS-UTILS**8.124.1. RHBA-2013:1714 — nfs-utils bug fix and enhancement update**

Updated nfs-utils packages that fix several bugs and add various enhancements are now available.

The nfs-utils packages provide a daemon for the kernel Network File System (NFS) server and related tools such as the mount.nfs, umount.nfs, and showmount.

Bug Fixes**BZ#889272**

When the "Background", "Foreground" or "timeo" options were set in multiple sections of the nfsmount.conf configuration file, each of those options were incorrectly present in the resulting parsed values. This update changes this behavior so that the first instance of either option overrides any previous ones.

In addition, configuration file options could have been incorrectly passed to the mount syscall from sections that were not relevant to the options that were being performed. The parser has been made more strict so that each option can appear at most four times: once for the system section, once for the server-specific section, once for the mount-specific section, and once for the command line mount options.

BZ#890146

Prior to this update, running "nfsstat -s -o rpc" command produced output with incorrect labels in a table header. With this update, the underlying source code has been adapted to make sure that all columns now have the correct name.

BZ#892235

Starting the nfs service resulted in the following output:

```
Stopping RPC idmapd:          [ OK ]
Starting RPC idmapd:         [ OK ]
```

Although the sequence of events of having to first stop and then start the RPC idmapd service was previously necessary, the current init scripts do not require this behavior. This has been corrected so that starting the nfs service now simply results in a single "Starting RPC idmapd" status display.

BZ#950324

When running sm-notify, specifying the "-v <ip_address>" or "-v <hostname>" option did not work correctly after the nfs-utils packages were updated to version 1.2.2, which was the first version that included support for IPv6. This update corrects the address handling logic so that specifying a hostname, IPv4 or IPv6 IP address with the '-v' option works as expected.

BZ#952560

The nfs(5) manual page contained incorrect information about the "retrans=n" option, by which specifies the number of times an NFS client will retry a request before it attempts a further recovery action. This information has been corrected and now specifies the number of attempts by protocol type. The man page correction for the "retrans=n" option is:

The number of times the NFS client retries a request before it attempts further recovery action. If the retrans option is not specified, the NFS client tries each request three times with mounts using UDP and two times with mounts using TCP.

Users of nfs-utils are advised to upgrade to these updated packages, which fix these bugs and add various enhancements.

8.125. NMAP

8.125.1. RHBA-2013:0881 — nmap bug fix update

Updated nmap packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The nmap packages provide a network exploration utility and a security scanner.

Bug Fixes

BZ#729045

Previously, the debuginfo file for the ncat utility was missing in the nmap debuginfo package. Consequently, debugging and analysis of unexpected terminations could not be done properly. This update ensure the missing file is present in the package, thus fixing this bug.

BZ#826601

In a previous version, the ncat utility failed to write its session data to an output file when the used protocol was UDP. This update provides a patch, which ensures that the data are properly written in the described scenario, thus fixing this bug.


Users of nmap are advised to upgrade to these updated packages, which fix these bugs.

8.126. NSS AND NSPR

8.126.1. RHBA-2013:1558 — nss and nspr bug fix and enhancement update

Updated nss and nspr packages that fix a number of bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Network Security Services (NSS) is a set of libraries designed to support the cross-platform development of security-enabled client and server applications. Netscape Portable Runtime (NSPR) provides platform independence for non-GUI operating system facilities.



NOTE

The nss family of packages, consisting of nss, nss-softokn, and nss-util, has been upgraded to the higher upstream versions, which provide a number of bug fixes and enhancements over the previous versions:

- The nss package has been upgraded to the upstream version 3.15.1. (BZ#918950, BZ#1002645)
- The nss-softokn package has been upgraded to the upstream version 3.14.3 (BZ#919172)
- The nss-util package has been upgraded to the upstream version 3.15.1 (BZ#919174, BZ#1002644)

The nspr package has been upgraded to upstream version 4.10, which provides a number of bug fixes and enhancements over the previous version. (BZ#919180, BZ#1002643)

Bug Fixes

BZ#702083

The PEM module imposed restrictions on client applications to use unique base file names upon which certificates were derived. Consequently, client applications certifications and keys with the same base name but different file paths failed to load because they were incorrectly deemed to be duplicates. The comparison algorithm has been modified and the PEM module now correctly determines uniqueness regardless of how users name their files.

BZ#882408

Due to differences in the upstream version of the nss package, an attempt to enable the unsupported **SSL PKCS#11 bypass** feature failed with a fatal error message. This behavior could break the

semantics of certain calls, thus breaking the Application Binary Interface (ABI) compatibility. With this update, the nss package has been modified to preserve the upstream behavior. As a result, an attempt to enable **SSL PKCS#11 bypass** no longer fails.

BZ#903017

Previously, there was a race condition in the certification code related to smart cards. Consequently, when Common Access Card (CAC) or Personal Identity Verification (PIV) smart cards certificates were viewed in the **Firefox** certificate manager, the **Firefox** web browser became unresponsive. The underlying source code has been modified to fix the race condition and **Firefox** no longer hangs in the described scenario.

BZ#905013

Due to errors in the Netscape Portable Runtime (NSPR) code responsible for thread synchronization, memory corruption sometimes occurred. Consequently, the web server daemon (**httpd**) sometimes terminated unexpectedly with a segmentation fault after making more than 1023 calls to the NSPR library. With this update, an improvement to the way NSPR frees previously allocated memory has been made and **httpd** no longer crashes in the described scenario.

BZ#918136

With the 3.14 upstream version of the nss package, support for certificate signatures using the MD5 hash algorithm in digital signatures has been disabled by default. However, certain websites still use MD5-based signatures and therefore an attempt to access such a website failed with an error. With this update, MD5 hash algorithm in digital signatures is supported again so that users can connect to the websites using this algorithm as expected.

BZ#976572

With this update, fixes to the implementation of Galois/Counter Mode (GCM) have been backported to the nss package since the upstream version 3.14.1. As a result, users can use GCM without any problems already documented and fixed in the upstream version.

BZ#977341

Previously, the output of the **certutil -H** command, which is a list of options and arguments used by the **certutil** utility, did not describe the **-F** option. This information has been added and the option is now properly described in the output of **certutil -H**.

BZ#988083

Previously, the **pkcs11n.h** header was missing certain constants to support the Transport Layer Security (TLS) 1.2 protocol. The constants have been added to the nss-util package and NSS now supports TLS 1.2 as expected.

BZ#990631

Previously, Network Security Service (NSS) reverted the permission rights for the **pkcs11.txt** file so that only the owner of the file could read it and write to it. This behavior overwrote other permissions specified by the user. Consequently, users were prevented from adding security modules to their own configuration using the system-wide security databases. This update provides a patch to fix this bug. As a result, NSS preserves the existing permissions for **pkcs11.txt** and users are now able to modify the NSS security module database.

BZ#1008534

Due to a bug in Network Security Services (NSS), the installation of the IPA (Identity, Policy, Audit) server terminated unexpectedly and an error was returned. This bug has been fixed with this update and installation of the IPA server now proceeds as expected.

BZ#1010224

The NSS **softoken** cryptographic module did not ensure whether the **freebl** library had been properly initialized before running its self test. Consequently, certain clients, such as the Lightweight Directory Access Protocol (LDAP) client, could initialize and finalize NSS. In such a case, **freebl** was cleaned up and unloaded. When the library was loaded again, an attempt to run the test terminated unexpectedly causing client failures such as Transport Layer Security (TLS) connection errors. This bug has been fixed and **softoken** now correctly initializes **freebl** before running self tests. As a result, the failures no longer occur in the described scenario.

Enhancements

BZ#960193, BZ#960208

Network Security Services's (NSS) own internal cryptographic module in Red Hat Enterprise Linux 6.5 now supports the NIST Suite B set of recommended algorithms for Elliptic curve cryptography (ECC).

Users of `nss` and `nspr` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. After installing this update, applications using NSS or NSPR must be restarted for this update to take effect.

8.127. NTP

8.127.1. RHBA-2013:1593 — ntp bug fix and enhancement update

Updated `ntp` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Network Time Protocol (NTP) is used to synchronize a computer's time with another reference time source. The `ntp` packages include the `ntpd` daemon and utilities used to query and configure `ntpd`.



NOTE

The `ntp` packages have been upgraded to upstream version 4.2.6p5, which provides a number of bug fixes and enhancements over the previous version. (BZ#654004)

Bug Fixes

BZ#673198

The `ntpdate` service did not wait for the NetworkManager service to configure the network before attempting to obtain the date and time update from the Internet. Consequently, `ntpdate` failed to set the system clock if the network was not configured. With this update, `ntpdate` attempts to obtain updates from the Internet in several increasing intervals if the initial attempt fails. The system clock is now set even when NetworkManager takes longer period of time to configure the network.

BZ#749530

The `ntp-keygen` utility always used the DES-CBC (Data Encryption Standard-Cipher Block Chaining)

encryption algorithm to encrypt private NTP keys. However, DES-CBC is not supported in FIPS mode. Therefore, ntp-keygen generated empty private keys when it was used on systems with FIPS mode enabled. To solve this problem, a new "-C" option has been added to ntp-keygen that allows for selection of an encryption algorithm for private key files. Private NTP keys are now generated as expected on systems with FIPS mode enabled.

BZ#830821

The ntpstat utility did not include the root delay in the "time correct to within" value so the real maximum errors could have been larger than values reported by ntpstat. The ntpstat utility has been fixed to include the root delay as expected and the "time correct to within" values displayed by the utility are now correct.

BZ#862983

When adding NTP servers that were provided by DHCP (using dhclient-script) to the ntp.conf file, the ntp script did not verify whether ntp.conf already contained these servers. This could result in duplicate NTP server entries in the configuration file. This update modifies the ntp script so that duplicate NTP server entries can no longer occur in the ntp.conf file.

BZ#973807

When ntpd was configured as a broadcast client, it did not update the broadcast socket upon change of the network configuration. Consequently, the broadcast client stopped working after the network service had been restarted. This update modifies ntpd to update the broadcast client socket after network interface update so the client continues working after the network service restart as expected.

Enhancements**BZ#623616, BZ#667524**

NTP now specifies four off-site NTP servers with the iburst configuration option in the default ntp.conf file, which results in faster initial time synchronization and improved reliability of the NTP service.

BZ#641800

Support for authentication using SHA1 symmetric keys has been added to NTP. SHA1 keys can be generated by the ntp-keygen utility and configured in the /etc/ntp/keys file on the client and server machines.

BZ#835155

Support for signed responses has been added to NTP. This is required when using Samba 4 as an Active Directory (AD) Domain Controller (DC).

BZ#918275

A new miscellaneous ntpd option, "interface", has been added. This option allows control of which network addresses ntpd opens and whether to drop incoming packets without processing or not. For more information on use of the "interface" option, refer to the ntp_misc(5) man page.

Users of ntp are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.128. NUMACTL

8.128.1. RHBA-2013:1712 — numactl bug fix update

Updated numactl packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The numactl packages provide a simple Non-Uniform Memory Access (NUMA) policy support and consist of the numactl program to run other programs with a specific NUMA policy and the libnuma library to do allocations in applications using the NUMA policy.

Bug Fixes

BZ#881779

Prior to this update, the "localalloc" option was not described clearly in the numactl(8) manual page, which could cause confusion. This update adds a clear description of "localalloc" to the numactl(8) manual page.

BZ#987507

Due to a bug in the numastat utility source code, output of the "numastat -m" command reported incorrect values of the amount of allocated static huge page memory. A patch has been applied to address this bug, and numastat now calculates huge page sizes properly.

Users of numactl are advised to upgrade to these updated packages, which fix these bugs.

8.129. NUMAD

8.129.1. RHBA-2013:1705 — numad bug fix and enhancement update

Updated numad packages that fix two bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The numad packages provide a daemon for NUMA (Non-Uniform Memory Architecture) systems, monitors NUMA characteristics and manages placement of processes and memory to minimize memory latency. The packages also provide an interface that can be used to query the numad daemon for the best manual placement of an application.

Bug Fixes

BZ#987563

When all CPUs were busy, the numad daemon was too reluctant to balance processes across nodes even though it would have resulted in significantly better application and system performance. With this update, numad is more aggressive about moving processes even when all CPUs are busy. As a result, overall system performance has improved significantly.

BZ#987559

Previously, it was not possible to set the system's hugepage "scan_sleep_millisecs" parameter. As a consequence, NUMA performance was damaged when process memory was migrated across nodes. The underlying code has been changed to accept the new "-H" option to specify "scan_sleep_millisecs", thus fixing the bug. The user can now set the value to fine-tune the numad's performance.

Enhancement

BZ#913546

A configuration file for the logrotate tool has been added to be fully supported by Red Hat Enterprise Linux 6.

All users of numad are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.130. OPENCRIPTOKI

8.130.1. RHBA-2013:1592 — opencryptoki bug fix and enhancement update

Updated opencryptoki packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The opencryptoki packages contain version 2.11 of the PKCS#11 API, implemented for IBM Cryptocards. This package includes support for the IBM 4758 Cryptographic CoProcessor (with the PKCS#11 firmware loaded), the IBM eServer Cryptographic Accelerator (FC 4960 on IBM eServer System p), the IBM Crypto Express2 (FC 0863 or FC 0870 on IBM System z), and the IBM CP Assist for Cryptographic Function (FC 3863 on IBM System z).

**NOTE**

The opencryptoki package has been upgraded to upstream version 2.4.3.1, which, compared to the previous version, provides support for the SHA-2 hash algorithms in the ICA token and adds fixes for the SHA-2- based certificates in the CCA token. (BZ#948349)

Users of opencryptoki are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.131. OPENCV

8.131.1. RHBA-2013:1118 — opencv bug fix update

Updated opencv packages that fix one bug are now available for Red Hat Enterprise Linux 6.

OpenCV is the open source computer vision library. It is a collection of C functions and C++ classes that implement Image Processing and Computer Vision algorithms.

Bug Fix**BZ#658060**

The OpenCVConfig.cmake file had different contents on 32-bit and 64-bit architecture and was installed under the /usr/share directory. Consequently, the opencv-devel package could not be installed in a multilib environment. With this update, the OpenCVConfig.cmake file has been moved to the /usr/lib(64) directory and the opencv-devel package can now be installed in a multilib environment.

Users of opencv are advised to upgrade to these updated packages, which fix this bug.

8.132. OPENHPI

8.132.1. RHBA-2013:1532 — openhpi bug fix update

Updated openhpi packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

OpenHPI provides an open source implementation of the Service Availability Forum (SAF) Hardware Platform Interface (HPI). HPI is an abstracted interface for managing computer hardware, typically chassis- and rack-based servers. HPI includes resource modeling; access to and control over sensor, control, watchdog, and inventory data associated with resources; abstracted System Event Log interfaces; hardware events and alarms; and a managed hot swap interface.

Bug Fixes

BZ#891626

Due to a bug in the `power_supply()` parsing routines, some returned strings could contain incorrectly displayed characters. Consequently, retrieving a serial or part number of a power supply unit (PSU) via the OpenHPI API resulted in strings containing these characters. This update ensures that proper serial and part numbers are returned for PSUs and the returned strings now only contain valid characters.

BZ#924852

Previously, code supporting certain RDR (Request Data with Reply) sensors was missing in OpenHPI. Consequently, after the extraction and reinsertion of an enclosure monitored via the Onboard Administrator (OA) SOAP plug-in, the following error messages were returned to the log file:

```
openhpid: ERROR: (oa_soap_sensor.c, 2005, RDR not present) openhpid: ERROR:
(oa_soap_fan_event.c, 279, processing the sensor event for sensor 24 has failed)
```

This bug has been fixed and no error messages are now logged after a component is extracted and reinserted.

BZ#948386

Under certain conditions, when using OpenHPI with the Onboard Administrator (OA) SOAP plug-in when an OA switch-over took place, HPI clients became unresponsive or the openhpi daemon failed to connect to the new active OA. Consequently, clients were unable to retrieve events and data. A series of patches has been provided to better account for OA failover situations, thus fixing this bug.

BZ#953515

Prior to this update, support for certain blade servers was missing in OpenHPI. Consequently, the OpenHPI daemon terminated unexpectedly with a segmentation fault at startup on these servers. A patch has been provided to add the missing support and the OpenHPI daemon no longer crashes in the described scenario.

BZ#953525

Due to missing support for certain thermal sensors, the `getBladeInfo()` function could terminate unexpectedly, causing the whole discovery process to fail. This update adds the support for these sensors and OpenHPI discovery now works as expected.

Users of openhpi are advised to upgrade to these updated packages, which fix these bugs.

8.133. OPENSCAP

8.133.1. RHBA-2013:1590 — openscap bug fix and enhancement update

Updated openscap packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The openscap packages provide OpenSCAP, which is a set of open source libraries for the integration of the Security Content Automation Protocol (SCAP). SCAP is a line of standards that provide a standard language for the expression of Computer Network Defense (CND) related information.



NOTE

The openscap packages have been upgraded to upstream version 0.9.12, which provides a number of bug fixes and enhancements over the previous version. This update adds support for the National Institute of Standards and Technology's (NIST) SCAP 1.2 standard, so that all content, such as the following, is correctly supported: the Red Hat Enterprise Linux 5 Security Technical Implementation Guide (STIG), The United States Government Configuration Baseline (USGCB), and Red Hat Security Advisory content. (BZ#[956763](#))

Bug Fix

BZ#[999903](#)

Previously, the oscap utility did not properly handle the process of object evaluation while querying the RPM database (RPMDB). RPMDB iterators created upon the query were not correctly removed if the process was aborted, which led to RPMDB corruption. With this update, the created RPMDB iterators are now removed correctly and process abortion no longer causes RPMDB corruption.

Users of openscap are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.134. OPENSSSH

8.134.1. RHSA-2013:1591 — Low: openssh security, bug fix, and enhancement update

Updated openssh packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE link(s) associated with each description below.

OpenSSH is OpenBSD's Secure Shell (SSH) protocol implementation. These packages include the core files necessary for the OpenSSH client and server.

Security Fix

CVE-2010-5107

The default OpenSSH configuration made it easy for remote attackers to exhaust unauthorized connection slots and prevent other users from being able to log in to a system. This flaw has been

addressed by enabling random early connection drops by setting `MaxStartups` to `10:30:100` by default. For more information, refer to the `sshd_config(5)` man page.

Bug Fixes

BZ#872169

An existing `/dev/log` socket is needed when logging using the `syslog` utility, which is not possible for all `chroot` environments based on the user's home directories. Previously, to fix this, a patch was applied to keep the `syslog` file descriptor open. However, the `syslog` library was changed and the used heuristic stopped working. As a consequence, the `sftp` commands were not logged in the `chroot` setup in the internal `sftp` subsystem. The patch has been adjusted to the new conditions and the `sftp` commands are logged in the `chroot` setup in the internal `sftp` subsystem.

BZ#880575

Previously, when the user attempted to use their own unprotected private key, the `ssh` utility displayed the following message:

```
It is recommended that your private key files are NOT accessible by others.
```

The key was subsequently rejected, which could have led to confusion as the behavior was inconsistent with the message. With this update, the message has been changed to:

```
It is required that your private key files are NOT accessible by others.
```

BZ#896561

The `ssh-agent` utility was unable to open more connections and could become unresponsive due to a race condition. The race condition has been fixed and `ssh-agent` no longer hangs in this scenario.

BZ#954094

If the `"bindpw"` option contained double quotes, it was not correctly parsed by the `ssh-ldap-helper` parser, and `ssh-ldap-helper` failed to bind to an LDAP server. With this update, `ssh-ldap-helper` parses the LDAP configuration files correctly.

BZ#955792

Prior to this update, non-ASCII characters have been replaced by their octal representations in banner messages in order to prevent terminal re-programming attacks. Consequently, banners containing UTF-8 strings were not correctly displayed in a client. With this update, banner messages are processed according to RFC 3454, control characters have been removed, and banners containing UTF-8 strings are now displayed correctly.

BZ#974096

Previously, if the `/tmp/` directory of the target user was polyinstantiated, no credentials cache was found on the remote machine after the Pluggable Authentication Module (PAM) session was initiated. As a consequence, Kerberos ticket forwarding did not work. With this update, the cache is re-created in a new `/tmp/` directory after the PAM session is initiated, and Kerberos ticket forwarding now works as expected.

BZ#993509

Previously, if the `sshd` daemon was configured to force the internal SFTP session, the daemon was

unable to properly handle requests for an interactive session. Consequently, sshd did not terminate SSH connections and SSH clients could become unresponsive. With this update, sshd has been modified to return an error message that the service allows SFTP connections only, and the SSH clients no longer hang in this scenario.

Enhancements

BZ#906872

This update adds support for certificate authentication of users and hosts using a new OpenSSH certificate format. Certificates contain a public key, identity information, and validity constraints, and are signed with a standard SSH public key using the ssh-keygen utility. Note that the version of ssh-keygen shipped with Red Hat Enterprise Linux 6 uses the "-Z" option for specifying the principals. For more information on this functionality, refer to the `/usr/share/doc/openssh-5.3p1/PROTOCOL.certkeys` file.

BZ#908038

This update adds support for PKCS#11 tokens. Now, OpenSSH clients are able to use smart cards for authentication.

BZ#951704

The `KexAlgorithms` configuration option has been added to client and server configuration in both the ssh utility and the sshd daemon. Specifying `KexAlgorithms` enables the user and the administrator to select key exchange methods and their order or preference.

BZ#969565

This update adds support for the SHA-2 Secure Hash Algorithm in the Hash-based Message Authentication Code (HMAC) to OpenSSH.

BZ#993577

The new Federal Information Processing Standard (FIPS) validation requires the random number generator (RNG) seed to have at least 112 bits of entropy instead of previous 80 bits. Therefore, the minimum value of the `SSH_USE_STRONG_RNG` environment variable has been increased to 14.

BZ#1001565

The new Federal Information Processing Standard (FIPS) validation requires the Power On Self Test (POST) to run in all cases when the FIPS module is installed. With this update, the POST self test is run on the SSH client and the SSH server if the `dracut-fips` package has been installed.

All openssh users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.135. OPENSSSL

8.135.1. RHBA-2013:1585 — openssl bug fix and enhancement update

Updated openssl packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The openssl packages provide a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols, as well as a full-strength general-purpose cryptography library.



NOTE

The openssl packages have been upgraded to upstream version 1.0.1e, which provides a number of bug fixes and enhancements over the previous version, including support for multiple new cryptographic algorithms and support for the new versions (1.1, 1.2) of the transport layer security (TLS) protocol. This update adds the following ciphers needed for transparent encryption and authentication support in GlusterFS: Cipher-based MAC (CMAC), XEX Tweakable Block Cipher with Ciphertext Stealing (AES-XTS), and Galois Counter Mode (AES-GCM). The following new additional algorithms are now supported: ECDH, ECDSA, and AES-CCM. (BZ#924250)

Bug Fixes

BZ#830109

Previously, an incorrect variable size was passed to the `getsockopt()` function. As a consequence, using the BIO (OpenSSL I/O) layer in datagram mode caused termination with a segmentation fault. More specifically, the openssl `s_client` command terminated unexpectedly on IBM System z with the `"-dtls1"` option enabled. After this update, a correctly-sized variable is used, and the datagram BIO functions no longer terminate with a segmentation fault on System z.

BZ#919404

Prior to this update, the `getaddrinfo()` function returned an error that was handled incorrectly in the openssl `s_server` command implementation. Consequently, the OpenSSL `s_server` did not work on IPv4-only systems. With this update, when `getaddrinfo()` fails on IPv6 addresses, the code has been modified to fall back to the IPv4 address lookup. As a result, the openssl `s_server` now correctly starts up on a computer with only IPv4 addresses configured.

Enhancements

BZ#818446

The Intel RDRAND instruction is now used, when available, to generate random numbers and has replaced the default OpenSSL random number generator. The instruction is not used when OpenSSL runs in FIPS mode.

BZ#929291

The performance of OpenSSL on current IBM PowerPC processors has been improved.

BZ#951690

The elliptic curve digital signature algorithm (ECDSA) and elliptic curve Diffie–Hellman (ECDH) algorithms are now enabled in OpenSSL. These algorithms support only elliptic curves listed in the national institute of standards and technology (NIST) Suite B specification.

BZ#951701

The new `"-trusted_first"` option has been added to OpenSSL. This enables preferring locally stored intermediate certificates instead of the intermediate certificates sent by the TLS server.

BZ#969562

Versions 1.1 and 1.2 of the transport layer security (TLS) protocol are now supported by the OpenSSL library.

BZ#969564

With this update, the "%{_prefix}" macro is used instead of the hardcoded /usr/ directory in the openssl.spec file when configuring OpenSSL before building.

BZ#987411

The next protocol negotiation (NPN) extension of the TLS protocol is now supported by OpenSSL. This extension allows for negotiation of the application protocol, which is used by the application, during the TLS handshake.

BZ#993584, BZ#999867

Due to the FIPS validation requirements, the FIPS Power-on self-tests (POST) always have to run when the FIPS module is installed. For libraries, this is ensured by running the self-tests from the dynamic library constructor function. If the dracut-fips package is installed, OpenSSL now treats it as an indicator that the OpenSSL FIPS module is installed and complete, and the self-tests run whenever the OpenSSL dynamic library is loaded.

Users of openssl are advised to upgrade to these updated packages, which fix these bugs and add these enhancements. For the update to take effect, all services linked to the OpenSSL library must be restarted, or the system rebooted.

8.135.2. RHBA-2013:1751 — openssl bug fix update

Updated openssl packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The openssl packages provide a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols, as well as a full-strength general-purpose cryptography library.

Bug Fixes**BZ#1025597**

Previously, the OpenSSL code incorrectly used RDRAND instruction when running on Cyrix CPU, which does not support it. Consequently, the applications that use the OpenSSL utility terminated unexpectedly on startup. The detection of CPU features on Cyrix CPU has been fixed, and the applications using OpenSSL no longer crash in the described scenario.

BZ#1025598

Prior to this update, the Transport Layer Security (TLS) client advertised support for some elliptic curves that are not supported by it. As a consequence, server could choose unsupported elliptic curve and client would not be able to communicate with the server over the TLS. With this update, OpenSSL TLS client advertises only the curves that are supported by it, and TLS communication with server (using also curves not supported by the Red Hat Enterprise Linux OpenSSL TLS client) can now be established.

Users of openssl are advised to upgrade to these updated packages, which fix these bugs. For the update to take effect, all services linked to the OpenSSL library must be restarted, or the system rebooted.

8.136. OPENSWAN**8.136.1. RHBA-2013:1718 — openswan bug fix and enhancement update**

Updated openswan packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Openswan is a free implementation of Internet Protocol Security (IPsec) and Internet Key Exchange (IKE). IPsec uses strong cryptography to provide both authentication and encryption services. These services allow you to build secure tunnels through untrusted networks.

Bug Fixes

BZ#771612

Previously, the "ipsec barf" command called the grep utility on the /var/log/lastlog file which caused the system to use significant amount of memory. After this update, "ipsec barf" uses the "lastlog -u user" command, which prevents the utility from using too much memory.

BZ#831669

According to the RFC 5996 standard, reserved fields must be ignored on receipt, irrespective of their value. Previously, however, the contents of the reserved fields was not being ignored on receipt for some payloads. Consequently, Openswan reported an error message and Internet Key Exchange (IKE) negotiation failed. With this update, Openswan has been modified to ignore the reserved fields and IKE negotiation succeeds regardless of the reserved field value.

BZ#831676

When a connection was configured in transport mode, Openswan did not pass information about traffic selectors to the NETKEY/XFRM IPsec kernel stack during the setup of security associations (SAs). Consequently, the information was not available in the output of the "ip xfrm state" command. With this update, Openswan correctly passes the traffic selectors information to the kernel when SAs are set up in transport mode.

BZ#846797

When a tunnel was established between two IPsec hosts, for example host1 and host2, utilizing Dead Peer Detection DPD, and if host2 went offline while host1 continued to transmit data, host1 continually queued multiple phase 2 requests after the DPD action. When host2 came back online, the stack of pending phase 2 requests was established, leaving a new IPsec Security Association (SA), and a large group of extra SA's that consumed system resources and eventually expired. This update ensures that Openswan has just a single pending phase 2 request during the time that host2 is down, and when host2 comes back up, only a single new IPsec SA is established, thus preventing this bug.

BZ#848132

When a tunnel was established between two IPsec hosts, for example host1 and host2, using the "dpdaction=restart" option, if host2 went offline and the Dead Peer Detection (DPD) was activated, the new phase1 replacement started retransmitting, but was subject to a limited amount of retries, even if the "keyingtries=%forever" option (which is default) was set. If host2 did not reconnect in time, the phase1 replacement expired and then the tunnel did not rekey until the old phase1 Security Association (SA) expired (in about 10 minutes by default). This meant that using the "dpdaction=restart" option only allowed a short window for the peer to reconnect. With this update, the phase1 replacement continues to try to rekey, thus avoiding the retransmission limit and timeout.

BZ#868986

Previously, certificates specified by names in "rightid" connection options containing a comma, were ignored and these connections were not authenticated due to an ID mismatch. With this update, Openswan now supports escaped commas inside the OID field in the "rightid" option.

BZ#881914

Previously, when certificates signed with the SHA2 digest algorithm were used for peer authentication, connection setup failed with the following error:

```
digest algorithm not supported
```

This bug has been fixed and Openswan now recognizes these certificates and sets up a connection correctly.

BZ#954249

The openswan package for Internet Protocol Security (IPsec) contains two diagnostic commands, "ipsec barf" and "ipsec look", that can cause the iptables kernel modules for NAT and IP connection tracking to be loaded. On very busy systems, loading such kernel modules can result in severely degraded performance or lead to a crash when the kernel runs out of resources. With this update, the diagnostic commands do not cause loading of the NAT and IP connection tracking modules. This update does not affect systems that already use IP connection tracking or NAT as the iptables and ip6tables services will already have loaded these kernel modules.

BZ#958969

Previously, when the IPsec daemon (pluto) attempted to verify the signature of a Certificate Revocation List (CRL), if the signature value began with a zero byte and had another zero as padding, the mpz() functions stripped out all leading zeros. This resulted in the Network Security Services (NSS) data input being one byte short and consequently failing verification when NSS compared its length to the modulus length. This update removes the conversions into arbitrary-precision arithmetic (bignum) objects and handles the leading zero by moving the pointer one position forward and reducing the length of the signature by 1. As a result, verification of CRLs now works as expected even with leading zeros in the signature.

BZ#960171

Previously, the order of the load_crls() and load_authcerts_from_nss() functions in the plutomain.c file was incorrect. As a consequence, when the IPsec daemon (pluto) attempted to load the Certificate Revocation Lists (CRLs) from the /etc/ipsec.d/crls/ directory during startup, loading failed because pluto checked for a loaded Certification Authority (CA) when there was none available. This update swaps the order of the aforementioned functions in the plutomain.c file, and now pluto no longer fails during startup and loads the CRLs successfully.

BZ#965014

Previously, the Openswan Internet Key Exchange version 2 (IKEv2) implementation did not set the "reserved" field to zero. As a consequence, Openswan did not pass the TAHI IKEv2 test. After this update, Openswan now sets the "reserved" field to zero and successfully passes the TAHI IKEv2 test.

BZ#975550

Previously, when an MD5 hash was used in the Internet Key Exchange version 2 (IKEv2) algorithm in Openswan to connect to another IPsec implementation, for example strongswan, occasionally the installed kernel security policy entry had a different "enc" or "auth" value than the corresponding values on the other side. As a consequence, a connection could not be established even though the Security Association (SA) was established correctly. After this update, these values are set correctly in Openswan and a connection can be established successfully.

BZ#985596

Previously, when in FIPS mode, Openswan did not allow the use of SHA2 algorithms. This update enables the use of SHA2 algorithms in FIPS mode.

BZ#994240

Initial support for passing traffic selectors to an XFRM IPsec stack for transport mode was incomplete and did not include the necessary work-arounds for NAT-traversal support. As a consequence, Openswan could not establish an L2TP connection with devices which use NAT-Traversal. After this update, the direction of IPsec Security Association (SA) is now passed to the `netlink_setup_sa()` function so that the client IP is substituted with the host IP and the selector works for NAT transport mode.

BZ#1002633

After this update, Openswan now uses `dracut-fips` to determine whether it should run in FIPS mode.

Enhancements**BZ#916743**

This update introduces a feature to control transmission delay timing for IPsec connections.

BZ#880004

With this update, Openswan now supports Internet Key Exchange (IKE) fragmentation. Openswan can now successfully connect to devices which support IKE fragmentation.

BZ#908476

Support for the Internet Key Exchange version 1 (IKEv1) INITIAL-CONTACT IPsec message, as defined in Section 4.6.3.3. of the RFC2407 specification, has been added to Openswan. This addresses an interoperability bug where a peer does not replace an existing IPsec Security Association (SA) with a newly negotiated one unless a Notification Payload message is present.

BZ#957400

The kernel module `aesni_intel` is now loaded by Openswan on startup. This update significantly improves the performance of Openswan on machines running Advanced Encryption Standard New Instructions (AES-NI).

BZ#959568

The default behavior of Openswan is to send NAT-Traversal keepalive packets. Disabling sending keepalive packets previously was a global option. After this update, the user can disable NAT-Traversal keepalive packet sending per connection.

Users of openswan are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.137. PACEMAKER

8.137.1. RHSA-2013:1635 — Low: pacemaker security, bug fix, and enhancement update

Updated pacemaker packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available from the CVE link associated with the description below.

Pacemaker is a high-availability cluster resource manager with a powerful policy engine.

Security Fix

CVE-2013-0281

A denial of service flaw was found in the way Pacemaker performed authentication and processing of remote connections in certain circumstances. When Pacemaker was configured to allow remote Cluster Information Base (CIB) configuration or resource management, a remote attacker could use this flaw to cause Pacemaker to block indefinitely (preventing it from serving other requests).



NOTE

The default Pacemaker configuration in Red Hat Enterprise Linux 6 has the remote CIB management functionality disabled.



NOTE

The pacemaker package has been upgraded to upstream version 1.1.10, which provides a number of bug fixes and enhancements over the previous version:

- * Pacemaker no longer assumes unknown cman nodes are safely stopped.
- * The core dump file now converts all exit codes into positive 'errno' values.
- * Pacemaker ensures a return to a stable state after too many fencing failures, and initiates a shutdown if a node claimed to be fenced is still active.
- * The `crm_error` tool adds the ability to list and print error symbols.
- * The `crm_resource` command allows individual resources to be reprobbed, and implements the `--ban` option for moving resources away from nodes. The `--clear` option has replaced the `--unmove` option. Also, `crm_resource` now supports OCF tracing when using the `--force` option.
- * The IPC mechanism restores the ability for members of the `haclient` group to connect to the cluster.
- * The Policy Engine daemon allows active nodes in the current membership to be fenced without quorum.
- * Policy Engine now suppresses meaningless IDs when displaying anonymous clone status, supports maintenance mode for a single node, and correctly handles the recovered resources before they are operated on.
- * XML configuration files are now checked for non-printing characters and replaced with their octal equivalent when exporting XML text. Also, a more reliable buffer allocation strategy has been implemented to prevent lockups.

(BZ#[987355](#))

Bug Fixes

BZ#902407

The "crm_resource --move" command was designed for atomic resources and could not handle resources on clones, masters, or slaves present on multiple nodes. Consequently, crm_resource could not obtain enough information to move a resource and did not perform any action. The "--ban" and "--clear" options have been added to allow the administrator to instruct the cluster unambiguously. Clone, master, and slave resources can now be navigated within the cluster as expected.

BZ#908450

The hacluster user account did not have a user identification (UID) or group identification (GID) number reserved on the system. Thus, UID and GID values were picked randomly during the installation process. The UID and GID number 189 was reserved for hacluster and is now used consistently for all installations.

BZ#913093

Certain clusters used node host names that did not match the output of the "uname -n" command. Thus, the default node name used by the crm_standby and crm_failcount commands was incorrect and caused the cluster to ignore the update by the administrator. The crm_node command is now used instead of the uname utility in helper scripts. As a result, the cluster behaves as expected.

BZ#951371

Due to incorrect return code handling, internal recovery logic of the crm_mon utility was not executed when a configuration updated failed to apply, leading to an assertion failure. Return codes are now checked correctly, and the recovery of an expected error state is now handled transparently.

BZ#996850

cman's automatic unfencing feature failed when combined with Pacemaker. Support for automated unfencing in Pacemaker has been added, and the unwanted behavior no longer occurs.

All pacemaker users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.138. PAM

8.138.1. RHEA-2013:1734 — pam enhancement update

Updated pam packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

Pluggable Authentication Modules (PAM) provide a system to set up authentication policies without the need to recompile programs to handle authentication.

Enhancement

BZ#976033

During TTY auditing, it is usually not necessary or even not desirable to log passwords that are being entered by the audited operator. This update adds an enhancement to the pam_tty_audit PAM module, so that passwords entered in the TTY console are logged only in case the "log_passwd"

option is used. As a result, passwords are no longer logged, unless the "log_passwd" option of pam_tty_audit is used. Note that this option is not available in kernel versions available prior to Red Hat Enterprise Linux 6.5.

Users of pam are advised to upgrade to these updated packages, which add this enhancement.

8.139. PAPI

8.139.1. RHBA-2013:1587 — papi bug fix and enhancement update

Updated papi packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

PAPI (Performance Application Programming Interface) is a software library that provides access to the processor's performance-monitoring hardware. This allows developers to track performance-related events, such as cache misses, instructions retired, and clock cycles, to better understand the performance issues of the software.



NOTE

The papi packages have been upgraded to upstream version 5.1.1, which provides a number of bug fixes and enhancements over the previous version, including support for Intel Xeon Processor E5-XXXX v2 architecture. (BZ#[831751](#))

Bug Fixes

BZ#[740909](#)

Due to missing dependencies in the makefile, a parallel rebuild of the PAPI library failed. With this update, new rules have been added to the makefile to address this problem. As a result, PAPI can be successfully rebuild in the described scenario.

BZ#[785258](#)

Previously, when Hyper-threading was enabled on the Intel Xeon Processor E5-XXXX node, the PAPI library could not configure the performance-monitoring hardware to count floating-point operations. This bug has been fixed and the aforementioned error no longer occurs.

BZ#[883475](#)

Due to an incorrect ldconfig setting in the papi.spec file, papi failed to be rebuilt from the srpm file when the process was executed by the root user. With this update, the underlying source code has been modified to fix this bug.

BZ#[883766](#)

Previously, the papi package failed to be built from the srpm file when a previous version of papi was installed. During the build, the new version of papi attempted to link to the libpfm.so file of the previously installed papi-devel package, which caused papi to terminate unexpectedly. With this update, a patch has been introduced to reorder the sequence of file linking during the build, so that the locally built files are used first. As a result, papi is built correctly with previous version installed.

Enhancements

BZ#726798, BZ#831751, BZ#947622

Support for the Intel Xeon Processor E5-XXXX and Intel Xeon Processor E5-XXXX architectures has been added to the PAPI library.

BZ#743648

Support for access to various energy and performance registers through PAPI has been added.

BZ#785975

With this update, several minor grammatical errors have been corrected in the PAPI interface.

BZ#866590

The papi-static subpackage has been added to provide the libraries for static linking.

BZ#910163

The papi-testsuite subpackage has been added to allow testing of papi.

All users of papi are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.140. PARTED

8.140.1. RHBA-2013:1627 — parted bug fix update

Updated parted packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The parted packages provide tools to create, destroy, resize, move, and copy hard disk partitions. The parted program can be used for creating space for new operating systems, reorganizing disk usage, and copying data to new hard disks.

Bug Fixes

BZ#851705

After removing a partition from a Device Mapper Multipath (DM-Multipath) device that contained the only partition, the parted utility failed to remove the DM-Multipath device as expected until system reboot. With this update, a bug in the libparted library has been fixed, so that parted now correctly removes DM-Multipath devices in this scenario.

BZ#869743

Previously, the parted utility was only able to handle loop devices with a size limited to an unsigned 32-bit integer. As a consequence, using parted with loop devices larger than 4 GiB failed. With this update, the above mentioned limit has been increased to an unsigned 64-bit integer, and loop devices up to 2^{64} of size are now supported.

BZ#631928

Previously, the parted utility was not able to handle Extended Address Volumes (EAV) Direct Access Storage Devices (DASD) that have more than 65535 cylinders. Consequently, EAV DASD drives could not be partitioned using parted, and installation on EAV DASD drives failed. This update adds support for EAV DASD devices with more than 65535 cylinders, and parted now handles them correctly.

Users of parted are advised to upgrade to these updated packages, which fix these bugs.

8.141. PCS

8.141.1. RHBA-2013:1633 — pcs bug fix and enhancement update

Updated pcs packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The pcs packages provide a command-line tool to configure and manage Pacemaker and Corosync tools.



NOTE

The pcs packages have been upgraded to upstream version 0.9.90, which provides a number of bug fixes and enhancements over the previous version, including improved stability and error checking. (BZ#[993115](#))

Bug Fixes

BZ#[901588](#)

Previously, the constraint rules IDs and the resource operation IDs were not displayed in the pcs utility. As a consequence, users were unable to remove the rules using pcs and had to use other tools or edit the Cluster Information Base (CIB) directly. This update adds the "--full" option which displays the IDs for resource operations and constraint rules, and users are now able to handle the rules using pcs.

BZ#[901607](#)

Previously, the pcs utility was unable to create constraints to promote or demote master and slave resources. Consequently, it was impossible to promote or demote master and slave resources using pcs. This update adds the ability for pcs to promote or demote master and slave resources using constraints.

BZ#[902450](#)

Previously, when the user created a resource with a monitor operation and then attempted to update the monitor operation, instead of updating the monitor operation a new operation was created. Consequently, it was impossible to update resource monitor operations in the pcs utility. This bug has been fixed and updating resource monitor operations in pcs now works as expected.

BZ#[902453](#)

The pcs utility showed no error message when the user selected a non-existent resource agent while creating a resource. Consequently, the resource did not start but the user was not notified of it. This update adds an error message, which can be overridden with the "--force" flag when necessary, and users are now properly notified in case of selecting a non-existent resource agent.

BZ#[902460](#)

Previously, the pcs utility contained no specific error message when the user submitted an unrecognized option. Consequently, pcs returned a traceback instead of an error message. With this update, pcs prints a specific error message in the described scenario.

BZ#[903712](#)

The "pcs config" command did not show resources configured as master and slave resources when the configuration of the Pacemaker cluster was reviewed. Consequently, users had to directly analyze the Cluster Information Base (CIB) to view the configured master and slave resources and their options. This update provides a patch to address this bug, so the "pcs config" command now shows all resources, including master and slave resources.

BZ#912496

Prior to this update, it was impossible to specify multiple resource operations when creating a resource in the pcs utility. Consequently, only resources with one operation could be created. The "pcs resource create" command has been updated and users are now able to create resources with multiple operations.

BZ#912498

Previously, when a resource had multiple operations, they were all displayed on the same line. Consequently, it was difficult to see all operations of a particular resource. With this update, each resource operation is displayed on its own line.

BZ#912528

Previously, the pcs utility did not contain a command for cleaning up resource failures. Consequently, users were unable to clean up the failed resources. This update adds the "pcs resource cleanup" command to solve this problem.

BZ#915248

Previously, the pcs utility did not support checking or resetting failure counts of a resource. As a consequence, users had to use other unsupported tools instead. With this update, users can reset and view failure counts for the specific resources using the "pcs resource failcount show" and "pcs resource failcount reset" commands.

BZ#916993

Prior to this update, the pcs utility had no manual page. As a consequence, users had to use the "pcs --help" command instead of viewing a manual page. The pcs(8) manual page has been added and users can now view full documentation for pcs using the "man pcs" command.

BZ#920767

Previously, the pcs utility did not contain support for managing or unmanaging resource groups. Consequently, users were allowed to manage and unmanage only individual resources. This update adds the necessary support and users can now manage and unmanage entire groups of resources.

BZ#998970

Previously, removing a group containing resources did not work correctly in the pcs utility. Consequently, the group was removed, but the resources remained. This bug has been fixed and the "pcs resource delete <groupname>" command now properly deletes both the group and the resources it contains.

Enhancement**BZ#887926**

This update adds support for editing cluster configuration with a standard editor to the pcs utility. This will save the user several steps, including saving the Cluster Information Base (CIB), editing it, and then pushing it back to the cluster. Users can edit and update the cluster configuration in one step

using the "pcs cluster edit" command.

Users of pcs are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.142. PERL

8.142.1. RHBA-2013:1534 — perl bug fix and enhancement update

Updated perl packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Perl is a high-level programming language that is commonly used for system administration utilities and web programming.



NOTE

The perl package has been upgraded to upstream version 2.021, which provides a number of bug fixes and enhancements over the previous version. Support for 64-bit ZIP archives has been improved. Especially, size of files bigger than 2³² bytes is now reported properly. (BZ#810469)

Bug Fixes

BZ#767608

Previously, referring to a named capturing group with non-matching name caused a memory leak. With this update, the underlying source code has been modified to avoid memory leaks in this scenario.

BZ#819042

When the `parse_file()` function from the `Pod::Man` or `Pod::Text` modules was executed without specifying the function output, `parse_file()` terminated. With this update, `parse_file()` has been modified to use standard output by default. As a result, `parse_file()` no longer fails with undefined output.

BZ#825713

Prior to this update, the `find2perl` utility incorrectly translated global expressions that contained the question mark ("?") character. Consequently, Perl code matched different expressions than the 'find' command-line utility. With this update, the global expression translator has been modified and `find2perl` now matches the same glob expressions as the 'find' utility does.

BZ#839788

Exiting scope of an object whose destructor method has been declared but not yet defined caused the Perl interpreter to terminate unexpectedly. This bug has been fixed and the interpreter now handles the undefined destructor methods as expected.

BZ#905482

When the XML-LibXSLT library was built without the `libgdm-devel` package installed on the system, it was unable to link to other libraries. With this update, the `glibc-devel`, `gdbm-devel`, and `db4-devel` packages have been added to the `perl-devel` list of run-time dependencies. As a result, it is now possible to build native Perl libraries without complications.

BZ#920132

While executing Perl code with the "format" option in a prototyped subroutine, the Perl interpreter terminated unexpectedly with a segmentation fault. With this update, various back-ported fixes have been added to the perl package. As a result, it is now possible to use formats in prototyped subroutines without complications.

BZ#973022

Prior to this update, the XML::Simple::XMLin() parser did not process input from the Getopt::Long::GetOptions() handler. Consequently, XML::Simple::XMLin() reported an unsupported method. With this update, Getopt::Long::GetOptions() has been modified to produce a simple string output that other Perl modules can read without complications.

BZ#991852

After installing a custom signal handler, the perl script attempted to access the thread-specific interpreter structure. This structure has already been disabled and Perl terminated with a segmentation fault. This bug has been fixed and Perl scripts no longer ask for the interpreter structure. As a result, Perl no longer crashes in the aforementioned scenario.

Enhancement**BZ#985791**

This update adds the CGI.pm module to the list of perl-core dependences. CGI.pm is now installed along with the perl-core package.

Users of perl are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.143. PERL-CGI-SESSION

8.143.1. RHBA-2013:0851 — perl-CGI-Session bug fix update

An updated perl-CGI-Session package that fixes one bug is now available for Red Hat Enterprise Linux 6.

CGI-Session is a Perl5 library that provides an easy, reliable and modular session management system across HTTP requests. Persistency is a key feature for such applications as shopping carts, login/authentication routines, and application that need to carry data across HTTP requests.

Bug Fix**BZ#657359**

Previously, several build-time dependencies were missing in the package. As a consequence, the package could not be rebuilt from source. This update adds the missing dependencies to the package, which can now be rebuilt from source as expected.

Users of perl-CGI-Session are advised to upgrade to this updated package, which fixes this bug.

8.144. PERL-CONFIG-GENERAL

8.144.1. RHBA-2013:1669 — perl-Config-General bug fix and enhancement update

Updated perl-Config-General packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Config::General Perl module is able to parse and write configuration files and presents an object-oriented access to their data. In addition, the Config::General module is 100% read-compatible with Apache configuration files, though it is also possible to create and access simple name/value configuration files. Certain enhancements such as here-documents, C-style comments and multi-line options are also available.

NOTE

The perl-Config-General package has been upgraded to upstream version 2.52, which provides a number of bug fixes and enhancements over the previous version. The most significant of these changes are the following:

- * When the Config::General module was supplied a directory instead of a configuration file, the module silently accepted it and returned an empty configuration hash. After this update, the module now calls the croak() function if the configuration file parameter is a directory and the "directory include" option is not turned on.
- * Parsing now accepts white spaces after the block closing character ">".
- * The save_file() function and a named block, whose second part starts with a slash character (/), now can be used to write valid configurations to a file correctly."
- * With this update, the sort() function is no longer used on arrays if the "-SaveSorted" option is turned off.
- * A character escaping bug has been fixed. Escaping the dollar "\$" or the backslash "\" characters is now handled correctly.
- * A tied hash now remains tied when saving to a file.
- * Preserving the single quotes during a variable interpolation has been corrected. An incrementor is now used to mark single quotes instead of using the rand() function.
- * Empty configuration values are no longer handed over to interpreting methods.
- * The "-Plug" parameter, which introduces plugin closures, has been added to the perl-Config-General packages.
- * The "-NoEscape" switch, which turns off escaping of every character, has been added to the perl-Config-General packages.
- * The parameters "-NormalizeOption", "-NormalizeBlock" and "-NormalizeValue", which take a subroutine reference and change a block, option or value accordingly, have been added to the perl-Config-General packages.
- * A new option "-AllowSingleQuoteInterpolation", which turns on interpolation for variables inside single quotes, has been added to the perl-Config-General packages.
- * A new option "-ForceArray" has been added to the perl-Config-General packages. When enabled, a single configuration value enclosed in square brackets "[]" will become an array forcefully.

With this upgraded perl-Config-General package, users can write and process their configuration files more efficiently. (BZ#658946)

Users of perl-Config-General are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.145. PERL-DATETIME

8.145.1. RHBA-2013:1566 — perl-DateTime bug fix update

Updated perl-DateTime packages that fix one bug are now available for Red Hat Enterprise Linux 6.

DateTime is a class for the representation of date/time combinations, and is part of the Perl DateTime project.

Bug Fix

BZ#978360

Previously, DateTime::Duration did not recognize the leap to 2012-07-01, which led to inaccurate computing time duration through the 2012-06-30T23:59:60 second. To fix this bug, a leap second appended in the end of 2012-06-30 has been added to perl-DateTime leap second database. Time arithmetic using Perl modules DateTime and DateTime::Duration now recognizes the leap from 2012-06-30 second correctly.

Users of perl-DateTime are advised to upgrade to these updated packages, which fix this bug.

8.146. PERL-MAKEFILE-PARSER

8.146.1. [RHBA-2013:0899](#) — perl-Makefile-Parser bug fix update

An updated perl-Makefile-Parser package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The perl-Makefile-Parser initial purpose is to provide basic support for the Makefile::GraphViz module, which is aimed to render the building process specified by a Makefile using the GraphViz library.

Bug Fix

BZ#657496

Previously, when the perl-Makefile-Parser source RPM package was built without the perl-Time-HiRes package installed, the building process failed executing the Makefile.PL script. With this update, build-time dependencies on the MDOM::Document::Gmake, MDOM::Util, and Time::HiRes Perl modules have been added to the RPM package and the perl-Makefile-Parser source package can now be built in minimal environment as expected.

Users of perl-Makefile-Parser are advised to upgrade to this updated package, which fixes this bug.

8.147. PERL-NET-DNS

8.147.1. [RHBA-2013:0785](#) — perl-Net-DNS bug fix update

Updated perl-Net-DNS packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The Perl packages provide the high-level programming language Perl, which is commonly used for system administration utilities and web programming.

Bug Fix

BZ#766357

Previously, dynamic update of an AAAA record caused the DNS module to return a FORMERR error on the prerequisite caused by the AAAA record creating the RDATA entry, even when the address

was never specified. Consequently, removing an AAAA record from a DNS zone failed. This update adds a check to ensure that required data are defined and removing AAAA records now works as expected.

Users of perl-Net-DNS are advised to upgrade to these updated packages, which fix this bug.

8.148. PERL-SOCKET6

8.148.1. [RHBA-2013:0777](#) — perl-`Socket6` bug fix update

Updated perl-`Socket6` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

This module supports the `getaddrinfo()` and `getnameinfo()` functions to intend to enable protocol independent programming. If the user's environment supports IPv6, IPv6 related defines such as `AF_INET6` are included.

Bug Fix

BZ#953873

When invoking a manual page for `Socket6`, no manual page was found. The build script has been fixed to convert POD (Plain Old Documentation) to manual pages. As a result, the "man `Socket6`" command correctly shows the `Socket6(3)` manual page for the `Socket6` Perl module as expected.

Users of perl-`Socket6` are advised to upgrade to these updated packages, which fix this bug.

8.149. PERL-TEST-MEMORY-CYCLE

8.149.1. [RHBA-2013:0859](#) — perl-`Test-Memory-Cycle` bug fix update

An updated perl-`Test-Memory-Cycle` package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The perl-`Test-Memory-Cycle` package provides possibility to search for circular references, which cannot be collected by the Perl's garbage collector.

Bug Fix

BZ#621089

Previously, the perl-`Test-Memory-Cycle` package was missing a build requirement. Consequently, the package could not be built. This update adds the perl-`CGI` package build requirement to the perl-`Test-Memory-Cycle.spec` file and perl-`Test-Memory-Cycle` can now be built as expected.

Users of perl-`Test-Memory-Cycle` are advised to upgrade to this updated package, which fixes this bug.

8.150. PERL-TEST-MOCKOBJECT

8.150.1. [RHBA-2013:0836](#) — perl-`Test-MockObject` bug fix update

Updated perl-Test-MockObject packages that fix one bug are now available for Red Hat Enterprise Linux 6.

Test::MockObject is a highly polymorphic testing object, capable of looking like all sorts of objects. This makes white-box testing much easier, as you can concentrate on what the code being tested sends to and receives from the mocked object, instead of worrying about making up your own data.

Bug Fix

BZ#661804

Building a perl-Test-MockObject source RPM package without an installed perl-CGI package failed on test execution. To fix this bug, build-time dependencies on the CGI, Test::Builder, and Test::More Perl modules have been declared in the RPM package. As a result, it is possible to rebuild the perl-Test-MockObject source RPM package in a minimal environment.

Users of perl-Test-MockObject are advised to upgrade to these updated packages, which fix this bug.

8.151. PERL-XML-DUMPER

8.151.1. [RHBA-2013:0838](#) — [perl-XML-Dumper bug fix update](#)

An updated perl-XML-Dumper package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The perl-XML-Dumper package provides the XML::Dumper module that allows converting Perl objects to XML format. XML::Dumper can also read XML data that was previously dumped by the module and convert it back to Perl objects. To ensure their correct behavior, the Perl objects should be converted and reconstituted in the same environment. When installed along with the Compress::Zlib module, XML::Dumper can also dump data to, and read data from, a compressed file with the ".xml.gz" extension.

Bug Fix

BZ#652833

The XML::Dumper module could not be used due to a missing dependency to the XML::Parser module. This update adds the required dependency to the perl-XML-Dumper spec file and XML::Dumper can now be used as expected.

Users of perl-XML-Dumper are advised to upgrade to this updated package, which fixes this bug.

8.152. PHP

8.152.1. [RHSA-2013:1615](#) — [Moderate: php security, bug fix, and enhancement update](#)

Updated php packages that fix three security issues, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

PHP is an HTML-embedded scripting language commonly used with the Apache HTTP Server.

Security Fixes

CVE-2006-7243

It was found that PHP did not properly handle file names with a NULL character. A remote attacker could possibly use this flaw to make a PHP script access unexpected files and bypass intended file system access restrictions.

CVE-2013-4248

A flaw was found in PHP's SSL client's hostname identity check when handling certificates that contain hostnames with NULL bytes. If an attacker was able to get a carefully crafted certificate signed by a trusted Certificate Authority, the attacker could use the certificate to conduct man-in-the-middle attacks to spoof SSL servers.

CVE-2013-1643

It was found that the PHP SOAP parser allowed the expansion of external XML entities during SOAP message parsing. A remote attacker could possibly use this flaw to read arbitrary files that are accessible to a PHP application using a SOAP extension.

Bug Fixes

BZ#892158, BZ#910466

Previously, when the `allow_call_time_pass_reference` setting was disabled, a virtual host on the Apache server could terminate with a segmentation fault when attempting to process certain PHP content. This bug has been fixed and virtual hosts no longer crash when `allow_call_time_pass_reference` is off.

BZ#947429

Prior to this update, if an error occurred during the operation of the `fclose()`, `file_put_contents()`, or `copy()` function, the function did not report it. This could have led to data loss. With this update, the aforementioned functions have been modified to properly report any errors.

BZ#969110

The internal buffer for the SQLSTATE error code can store maximum of 5 characters. Previously, when certain calls exceeded this limit, a buffer overflow occurred. With this update, messages longer than 5 characters are automatically replaced with the default "HY000" string, thus preventing the overflow.

Enhancement

BZ#953814

This update adds the following rpm macros to the php package: `%__php`, `%php_inidir`, `%php_inclidir`.

Users of php are advised to upgrade to these updated packages, which fix these bugs and add this enhancement. After installing the updated packages, the `httpd` daemon must be restarted for the update to take effect.

8.153. PIRANHA

8.153.1. RHBA-2013:1618 — piranha bug fix update

Updated piranha packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

Piranha provides high-availability and load-balancing services for Red Hat Enterprise Linux. The piranha packages contain various tools to administer and configure the Linux Virtual Server (LVS), as well as the heartbeat and failover components. LVS is a dynamically-adjusted kernel routing mechanism that provides load balancing, primarily for Web and FTP servers.

Bug Fixes

BZ#903711

Previously, the `lvsd` daemon did not properly activate the "sorry server" fallback service when all real servers were unavailable. Consequently, incoming traffic for a virtual service with no available real servers was not directed to "sorry server". This bug has been fixed and the `lvsd` daemon now properly activates "sorry server" when no real servers are available.

BZ#980169

In certain cases, most often caused by brief network outages, high latency between directors, or aggressive keepalive and deadtime settings, the `lvsd` daemon did not properly terminate when signaled by the pulse daemon. Consequently, `lvsd` ran on both master and backup directories having multiple virtual IPs (VIPs). With this update, `lvsd` has been modified to correctly catch and handle all signals from the pulse daemon. As a result, the redundant VIP is now properly removed.

Users of piranha are advised to upgrade to these updated packages, which fix these bugs.

8.154. 389-DS-BASE

8.154.1. RHBA-2013:1653 — 389-ds-base bug fix and enhancement update

Updated 389-ds-base packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The 389 Directory Server is an LDAPv3 compliant server. The base packages include the Lightweight Directory Access Protocol (LDAP) server and command-line utilities for server administration.

Bug Fixes

BZ#830334

Due to an incorrect interpretation of the error code, the Directory Server considered an invalid chaining configuration setting as the disk full error and terminated unexpectedly. Now, a more appropriate error code is used and the server no longer shuts down when invalid chaining configuration settings are specified.

BZ#905825

After the upgrade from Red Hat Enterprise Linux 6.3 to version 6.4, the upgrade script did not update the schema file for the `PamConfig` object class. Consequently, new features for PAM (Pluggable Authentication Module), such as configuration of multiple instances and `pamFilter` attribute, could not be used because of the schema violation. With this update, the upgrade script updates the schema file for the `PamConfig` object class as expected. As a result, the new features now function properly.

BZ#906005

Previously, the **valgrind** test suite reported recurring memory leaks in the **modify_update_last_modified_attr()** function. The size of these leaks averaged between 60-80 bytes per modify call, which could cause problems in environments with frequent modify operations. With this update, memory leaks no longer occur in the **modify_update_last_modified_attr()** function.

BZ#906583

Under certain circumstances, the **Directory Server** (DS) was not able to replace multi-valued attributes for new values that differed from the old ones only in the letter case. Consequently, a code 20 error message was displayed:

```
Type or value exists
```

With this update, **DS** has been modified to correctly process modification requests, and the letter case of attribute values can now be changed without complications.

BZ#907985

Under certain circumstances, the **DNA** (Distributed Numeric Assignment) plug-in logged messages with the **DB_LOCK_DEADLOCK** error code when attempting to create an entry with a `uidNumber` attribute. This bug has been fixed and **DNA** now handles this case properly and errors are no longer logged in the aforementioned scenario.

BZ#908861

The **Posix Winsync** plug-in was unnecessarily calling the internal **modify()** function. This internal **modify()** call failed and logged the following message:

```
slapi_modify_internal_set_pb: NULL parameter
```

With this update, **Posix Winsync** has been fixed and no longer calls **modify()**. As a result, the aforementioned message is no longer logged.

BZ#910581

Under certain circumstances, the `/etc/dirsrv/slapd-dstet-mkubik/dse.ldif` file was written with 0 bytes after a server termination or when the system was powered off. Consequently, after the system restart, the DS or IdM system sometimes did not start, leading to production server outages. The server mechanism by which `dse.ldif` is written has been modified, and server outages no longer occur in the described case.

BZ#913215

Prior to this update, while trying to remove a tombstone entry, the **ns-slapd** daemon terminated unexpectedly with a segmentation fault. This bug has been fixed and removal of tombstone entries no longer causes **ns-slapd** to crash.

BZ#921937

Previously, the **schema-reload** plug-in was not thread-safe. Consequently, executing the **schema-reload.pl** script under a heavy load could have caused the **ns-slapd** process to terminate unexpectedly with a segmentation fault. With this update, **schema-reload** has been modified to be thread-safe, and **schema-reload.pl** can be now executed along with other LDAP operations without complications.

BZ#923407

Due to an incorrect lock timing in the **DNA** (Distributed Numeric Assignment) plug-in, a deadlock occurred when **DNA** operation was executed along with other plug-ins. This update moves the release timing of the problematic lock, and **DNA** no longer causes the deadlock in the aforementioned scenario.

BZ#923502

Under certain circumstances, an out of scope local variable caused the **modrdrn** operation to terminate unexpectedly with a segmentation fault. This update modifies the declaration of the local variable so it does not get out of scope. As a result, **modrdrn** operations no longer crash.

BZ#923503

Previously, the **cleanallruv** task with the **replica-force-cleaning** option enabled did not remove all configuration attributes. Consequently, the task was initiated each time the server was restarted. With this update, the **cleanallruv** search mechanism has been modified, and **cleanallruv** no longer restarts when the server is restarted.

BZ#923504

Due to a bug in the **Ac1** plug-in, when using the **getEffectiveRights** request on a non-existing entry, a NULL pointer dereference could have occurred. Consequently, the server terminated unexpectedly with a segmentation fault. With this update, **Ac1** has been modified to check for NULL entry pointers. As a result, the server no longer crashes and an appropriate error message is now displayed when using **getEffectiveRights** request on a non-existing entry.

BZ#923909

Due to an insufficient size of the default **sasl_io** buffer, SASL connections could have been refused by the server. With this update, the buffer size has been increased to 65,536 bytes. Moreover, users can increase this value with the **nsslapd-sasl-max-buffer-size** setting. As a result, SASL connections are now accepted without complications.

BZ#947583

Previously, the code responsible for replication conflict resolution in the 389-ds-base package did not work correctly in several cases, such as conflict DN generation, retrieving deleted parent entry, and examining the scope of a deleted entry. Consequently, an intermediate node entry with positive child count but without children could have been created. The server then refused to remove such an entry. This update fixes the replication conflict resolution code, thus preventing the incorrect node entry creation.

BZ#951616

Previously, if a group on the Active Directory contained a member that was in a container of not-synchronized type, synchronizing the group with the LDAP server was unsuccessful. Consequently, the valid members were not synchronized. With this update, the entries in such containers are omitted and the synchronization is now successful in the described case.

BZ#953052

Prior to this update, certain schema definitions in the 389-ds-base package did not comply with the LDAP RFC 2252 standard. Consequently, problems with LDAP clients could have occurred. With this update, these schema definitions have been corrected to be compliant with LDAP RFC 2252.

BZ#957305

Under a very high load of hundreds of simultaneous connections and operations, the **Directory Server** could have encountered a race condition in the connection handling code. Consequently, the server terminated unexpectedly with a segmentation fault. With this update, code that updates the connection objects has been moved into the connection **mutex** object. As a result, **Directory Server** does not crash under high loads.

BZ#957864

Prior to this update, the **Simple Paged Results** control did not support an asynchronous search. Consequently, if the **Directory Server** received large number of asynchronous search requests, some of the requests terminated with error 53:

```
LDAP_UNWILLING_TO_PERFORM
```

With this update, asynchronous search support has been implemented into **Simple Paged Results**. As a result, **Directory Server** safely handles intensive asynchronous search requests.

BZ#958522

Previously, when loading an entry from a database, the `str2entry_dupcheck()` function was called instead of the more appropriate `str2entry_fast()` function. This behavior has been changed and `str2entry_fast()` is now called in the described scenario.

BZ#962885

The upgrade of Red Hat Enterprise Linux Identity Management server changed the value of the `nsslapd-port` variable to "0" for security reasons. The `nsslapd-port` is also used to construct the RUV (Replica Update Vector) used by replication. Previously, if the replication startup code found a zero `nsslapd-port`, it removed the RUV. Consequently, replication became unresponsive. With this update, RUV is no longer removed in the aforementioned scenario, thus preventing the replication hang.

BZ#963234

Previously, an empty control list was not handled properly by the **Directory Server**. Consequently, a LDAP protocol error was returned. With this update, **Directory Server** has been modified to handle sequences of zero length correctly, thus preventing the error.

BZ#966781

When there was a request for a new LDAP connection at the same time as a request for a new LDAPS or LDAPi connection, the **Directory Server** processed only the LDAP request. With this update, **Directory Server** has been modified to process all listener requests at the same time.

BZ#968383

Prior to this update, an incorrect error code (`err=0`) was returned when creating an invalid external SASL bind. With this update, a proper error code (`err=48`) is returned in the aforementioned scenario.

BZ#968503

When the **Directory Server** (DS) encountered an error while it processed a **startTLS** request, the server attempted to write a response back to the client. Consequently, DS became unresponsive. With this update, DS has been modified to correctly processes **startTLS** requests even in case of network errors. As a result, DS no longer hangs in the aforementioned scenario.

BZ#969210

Previously, the size of the `backlog` parameter of the `listen()` function was set to "128".

Consequently, if the server processed a large amount of simultaneous connection requests, the server could have dropped connection requests due to exceeded **backlog** size. With this update, a **nsslapd-listen-backlog-size** attribute has been added to allow the **backlog** size to be changed.

BZ#970995

Previously, the disk monitoring feature of the **Directory Server** did not function properly. If logging functionality was set to "critical" and logging was disabled, the rotated logs were deleted. If the attribute **nsslapd-errorlog-level** was explicitly set to any value, even zero, the disk monitoring feature did not stop the **Directory Server** as expected. This update corrects the settings of the disk monitoring feature and the server shuts down when the critical threshold is reached.

BZ#971033

Prior to this update, the **connections** attribute that stores the number of currently connected clients was incorrectly incremented twice, both by the **disconnect_server_nomutex()** and **connection_reset()** function. Consequently, the attribute contained incorrect values. This bug has been fixed and **connections** now store the correct number of connected clients.

BZ#972976

When the **Directory Server** (DS) used both the replication and the **DNA** plug-in, and the client sent a sequence of ADD or DELETE requests for the same entry, DS returned the following message:

```
modify_switch_entries failed
```

This bug has been fixed, and the aforementioned message is no longer returned.

BZ#973583

The internal **password** attribute is not preserved after the **Directory Server** (DS) restart. Previously, an attempt to delete the **password** after restarting DS, caused DS to terminate unexpectedly. With this update, DS has been modified to check if the **password** attribute exists, and if no, to skip the deletion. As a result, DS no longer crashes in the described case.

BZ#974361

Prior to this update, when using the **account policy** plug-in to configure policies for individual users based on the createTimestamp attribute, the createTimestamp was overwritten after the consequent binding. Consequently, **account policy** failed to lock the user. With this update, createTimestamp is no longer modified after successful binding and **account policy** now locks users as expected.

BZ#974719

Under certain circumstances, an inconsistent behavior of the modrdn operation when processing a tombstone entry caused the **Directory Server** (DS) to terminate unexpectedly. With this update, DS has been modified to correctly process tombstones with modrdn, thus preventing the crash.

BZ#974875

Prior to this update, when an attribute was configured to be encrypted, the on-line import failed to encrypt this attribute on a server. This update allows encryption on the consumer side, during an on-line import, thus fixing this bug.

BZ#975243

Previously, after removing the createTimestamp attribute from the account policy, this attribute was still applied by the Directory Server (DS). This bug has been fixed, and createTimestamp can now be effectively removed from the DS account policy.

BZ#975250, BZ#979169

Previously, with a mix of concurrent search, update, and replication operations a deadlock could have occurred between the changelog readers, writers, and main database writers. Consequently, the update operations failed. With this update, a new *nsslapd-db-deadlock-policy* configuration parameter has been introduced. The default value of this parameter is set to **9**, which terminates the last locker in case of a deadlock. After changing this value to **6**, the locker with the fewest write locks is terminated, which is advised for users who encounter frequent deadlocks.

BZ#976546

Prior to this update, if certain requested attributes were skipped during a search, the returned attribute names and values were sometimes transformed to upper case. This update removes attributes that are not authorized from the requested attributes set, so that the names of returned attributes or values are preserved in the correct form.

BZ#979435

Previously, after modifying a single-valued attribute in a multi-master replication environment, this change was not replicated to other servers. With this update, code that handles replication updates has been changed. As a result, the modify operations on single-valued attributes are replicated correctly.

BZ#982325

Previously, setting the "nsslapd-disk-monitoring-threshold" attribute with the ldapmodify utility to a large value worked as expected; however, due to a bug in the ldapsearch utility, the threshold value was displayed as a negative number. This update corrects the bug in ldapsearch and correct threshold values are now displayed.

BZ#983091

Previously, the Directory Server (DS) was not properly freeing the memory used by old connections. Consequently, when opening and closing hundreds of connections per minute for a long period of time, a memory leak occurred. With this update, DS has been modified to release the memory used by old connections as expected. As a result, the memory leak no longer occurs in the aforementioned scenario.

BZ#986131

Due to the USN (Update Sequence Number) configuration, the initial value of the lastusn variable in the rootdse directory was displayed as "18446744073709551615" instead of expected "-1". This update adds a special treatment for initial lastusn. As a result, this value is set to "-1" as expected. If a negative value is found in the USN index file, it is reset to the initial value.

BZ#986424

With this update, several minor coding errors have been corrected to prevent possible memory leaks and stability issues.

BZ#986857

If logging functionality was not set to "critical", the mount point for the logs directory was incorrectly skipped during the disk space check. The processing of configuration settings has been fixed and the log directory is no longer skipped.

BZ#987703

Previously, memory leaks occurred when using the `set_krb5_creds()` function for the replication transport or bind. The underlying source code has been modified and the memory leaks no longer occur.

BZ#988562

When multiple clients were connected to the Directory Server (DS), each of them adding and deleting users, the server deadlock could have occurred. With this update, a patch has been introduced to prevent the deadlock.

BZ#989692

When a server-side sorting request was evaluated, the "sort type" parameter was registered only from the first attribute in the request and the following attributes were ignored even if having different "sort type" values. Consequently, the sorting operation was performed incorrectly. With this update, Directory Server has been modified so that the server-side sorting resets "sort type" for each sort attribute in the request. As a result, the sorting is now handled correctly.

BZ#1002260

Due to a schema error, the Directory Server (DS) failed to start after the system upgrade. This bug has been fixed, and DS now works correctly in the described case.

BZ#1006846

If a replication was configured before initializing the sub backend, the temporary sub suffix was not updated with the real sub suffix entry. Consequently, the server search failed to return entries under the sub suffix. With this update, when a real sub suffix is added, the temporary entry ID in the `entryrdn` index is replaced with the real entry ID. As a result, search successfully returns sub suffix entries.

BZ#1007452

With certain specific values of the `nsDS5ReplicaName` variable, the replication could have become corrupted. With this update, all replica names are handled correctly.

BZ#1008013

In certain cases, the Directory Server became unresponsive when processing multiple outgoing and incoming operations using the TLS or SSL protocol. The underlying source code has been modified and the server no longer hangs in this scenario.

BZ#1013735

Previously, if the Directory Server (DS) worked with replicas that did not support the `CLEANALLRUV` task, running this task made DS unresponsive. With this update, DS has been modified to skip replicas that do not support `CLEANALLRUV`, thus fixing this bug.

BZ#1016038

Previously, when checking an Active Directory (AD) entry was a subject of synchronization, just the direct child of the target was checked. Consequently, AD entries which were in a deeper level were not synchronized to the Directory Server. This bug has been fixed, and child directories of the target

are now synchronized at and all levels.

Users of 389-ds-base are advised to upgrade to these updated packages, which fix these bugs.

8.155. PKI-CORE

8.155.1. RHBA-2013:1682 — pki-core bug fix update

Updated pki-core packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

Red Hat Certificate System is an enterprise software system designed to manage enterprise public key infrastructure (PKI) deployments. PKI Core contains fundamental packages required by Red Hat Certificate System, which comprise the Certificate Authority (CA) subsystem.

Note: The Certificate Authority component provided by this advisory cannot be used as a standalone server. It is installed and operates as a part of Identity Management (the IPA component) in Red Hat Enterprise Linux.

Bug Fixes

BZ#887305

Previously, the `/var/run/pki/ca/` directory was assigned an incorrect SELinux context after the installation of the pki-ca package. With this update, the `restorecon` command is applied on `/var/run/pki/ca/` during the post-installation process. As a result, this directory is now labeled with the correct SELinux context.

BZ#895702, BZ#999055

Prior to this update, when the pki-ca daemon was restarted on the Red Hat Enterprise Linux 6.4 Identity Management server, AVC denials were reported. With this update, pki-ca has been modified and AVC denials are no longer reported in the aforementioned scenario.

BZ#998715

The pki-selinux package sets the file context for certain default paths, so that the context need not be set when Red Hat Certificate System instances are created. Prior to this update, when pki-selinux was installed, unnecessary warning messages were displayed if these paths did not yet exist. These messages are now suppressed.

All users of pki-core are advised to upgrade to these updated packages, which fix these bugs.

8.156. POLICYCOREUTILS

8.156.1. RHBA-2013:1608 — policycoreutils bug fix and enhancement update

Updated policycoreutils packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The policycoreutils packages contain the core utilities that are required for the basic operation of a Security-Enhanced Linux (SELinux) system and its policies.

Bug Fixes

BZ#860506

Previously, several semanage command-line options did not work as expected. With this update, these options have been corrected and now work properly.

BZ#868218

With this update, the semanage man page has been updated to be consistent with information contained in the semanage help page.

BZ#886059

Due to a bug in the polycoreutils package, installation of the ipa-server-selinux package failed. This bug has been fixed and ipa-server-selinux can now be installed without complications.

BZ#913175

Prior to this update, the sandbox utility did not accept symbolic links specified in the `/etc/sysconfig/sandbox` file. Consequently, executing the "sandbox -M" command failed with a "No such file or directory" message. The underlying source code has been modified and symlinks can now be configured in `/etc/sysconfig/sandbox` without complications.

BZ#916727

Previously, the fixfiles script did not recognize a change in the regular expression describing the `/sbin/ip6?tables-multi*` files. Consequently, these files were labeled incorrectly after updating the system with the yum update command. With this update, fixfiles has been corrected to accept the change of regular expression.

BZ#918460

Prior to this update, after executing the "semanage boolean -m" command, a traceback was returned. This bug has been fixed and tracebacks are no longer displayed in the aforementioned scenario.

BZ#928320, BZ#947504

Previously, the semanage utility did not allow to set the "none" context on directories and files. Consequently, the following message was displayed when attempting to do so:

```
/usr/sbin/semanage: Type none is invalid, must be a file or device type
```

This bug has been fixed, and it is now possible to set the "none" context without complications.

BZ#967728

Prior to this update, the "-o" option of the audit2allow command overwrote the contents of the output file instead of just appending the new content. This bug has been fixed, and "audit2allow -o" now works as expected.

BZ#984484

After attempting to enable a SELinux boolean that did not exist, a brief error message was generated. This update modifies this message to be more informative.

BZ#998974

After attempting to permanently change a boolean that did not exist, an incorrect error message was generated. This message has been modified to provide correct information about the cause of the problem.

Enhancement

BZ#916734

This update adds a possibility to exclude selected files when running the restorecon utility, which can significantly reduce execution times.

Users of polycoreutils are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.157. POWERTOP

8.157.1. RHBA-2013:1575 — powertop bug fix and enhancement update

Updated powertop packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

PowerTOP is a tool to detect all the software components that make a computer consume more than necessary power when idle. PowerTOP can be used to reduce power usage by running various commands on the system.



NOTE

The powertop package has been upgraded to upstream version 2.3, which provides a number of bug fixes and enhancements over the previous version including corrected handling of arbitrary interface names. Moreover, checks for usability of the ondemand governor have been added. Also, several changes have been made to enhance user experience and to simplify and improve power management profiling capabilities. (BZ#682378, BZ#697273, BZ#829800)

Bug Fix

BZ#998021

The default soft limit for per-process open file descriptors is 1024, and the default hard limit for per-process file descriptors is 4096. By using the performance counter subsystem, the PowerTOP tool could exceed the limit on complex systems. Consequently, an error message about missing kernel support for perf was displayed. This update adds a fix that temporarily increases both the soft and hard file descriptor limits for the current process to the kernel limit. If the kernel limit is still insufficient, PowerTOP now displays an error message indicating that the file descriptor limits should be manually increased.

Users of powertop are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.158. PYKICKSTART

8.158.1. RHBA-2013:1629 — pykickstart bug fix and enhancement update

An updated pykickstart package that fixes several bugs and adds one enhancement is now available for Red Hat Enterprise Linux 6.

The pykickstart package contains a python library for manipulating kickstart files.

Bug Fixes

BZ#886010

When combining the autopart command with other kickstart partitioning commands in one kickstart file, installation proceeded unexpectedly with errors. The underlying source code has been modified so that the user is notified and installation is aborted with a parse error if the kickstart file contains invalid partitioning.

BZ#924579

When a kickstart file specified two logical volumes with the same name, installation failed with this inappropriate error message:

```
AttributeError: 'LogVolData' object has no attribute 'device'
```

With this update, duplicate names are detected correctly and the appropriate error message is displayed.

BZ#966183

When running a kickstart file that contained the "network" command with the "--ipv6" option specified, installation could terminate unexpectedly with the following message:

```
TypeError: not all arguments converted during string formatting
```

This update applies a patch to fix this bug and kickstart files work as expected in the described scenario.

Enhancement

BZ#978252

This enhancement adds the ability to specify the "--ipv6gateway" option for the "network" command in kickstart files. As a result, both IPv4 and IPv6 default gateways can be specified for a network device configuration using the "network" command.

Users of pykickstart are advised to upgrade to this updated package, which fixes these bugs and adds this enhancement.

8.159. PYPARTED

8.159.1. RHBA-2013:1616 — pyparted bug fix update

Updated pyparted packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The pyparted packages contain Python bindings for the libparted library. They are primarily used by the Red Hat Enterprise Linux installation software.

Bug Fix

BZ#896024

Due to a bug in the underlying source code, an attempt to run the parted.version() function caused a system error to be returned. This bug has been fixed and parted.version() can now be executed as expected.

Users of `pyparted` are advised to upgrade to these updated packages, which fix this bug.

8.160. PYTHON

8.160.1. [RHSA-2013:1582](#) — Moderate: python security, bug fix, and enhancement update

Updated python packages that fix one security issue, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

Python is an interpreted, interactive, object-oriented programming language.

Security Fix

[CVE-2013-4238](#)

A flaw was found in the way the Python SSL module handled X.509 certificate fields that contain a NULL byte. An attacker could potentially exploit this flaw to conduct man-in-the-middle attacks to spoof SSL servers. Note that to exploit this issue, an attacker would need to obtain a carefully crafted certificate signed by an authority that the client trusts.

Bug Fixes

[BZ#521898](#)

Previously, several Python executables from the `python-tools` subpackage started with the `#!/usr/bin/env python` shebang. This made it harder to install and use alternative Python versions. With this update, the first line of these executables has been replaced with `#!/usr/bin/python` that explicitly refers to the system version of Python. As a result, a user-preferred version of Python can now be used without complications

[BZ#841937](#)

Prior to this update, the `sqlite3.Cursor.lastrowid` object did not accept an insert statement specified in the Turkish locale. Consequently, when installing Red Hat Enterprise Linux 6 with the graphical installer, selecting "Turkish" as the install language led to an installation failure. With this update, `sqlite3.Cursor.lastrowid` has been fixed and installation no longer fails under the Turkish locale.

[BZ#845802](#)

Previously, the `SysLogHandler` class inserted a UTF-8 byte order mark (BOM) into log messages. Consequently, these messages were evaluated as having the emergency priority level and were logged to all user consoles. With this update, `SysLogHandler` no longer appends a BOM to log messages, and messages are now assigned correct priority levels.

[BZ#893034](#)

Previously, the `random.py` script failed to import the `random` module when the `/dev/urandom` file did not exist on the system. This led subsequent programs, such as `Yum`, to terminate unexpectedly. This bug has been fixed, and `random.py` now works as expected even without `/dev/urandom`.

BZ#919163

The **WatchedFileHandler** class was sensitive to a race condition, which led to occasional errors. Consequently, rotating to a new log file failed. **WatchedFileHandler** has been fixed and the log rotation now works as expected.

BZ#928390

Prior to this update, Python did not read Alternative Subject Names from certain Secure Sockets Layer (SSL) certificates. Consequently, a false authentication failure could have occurred when checking the certificate host name. This update fixes the handling of Alternative Subject Names and false authentication errors no longer occur.

BZ#948025

Previously, the **SocketServer** module did not handle the system call interruption properly. This caused certain HTTP servers to terminate unexpectedly. With this update, **SocketServer** has been modified to handle the interruption and servers no longer crash in the aforementioned scenario.

BZ#958868

Passing the **timeout=None** argument to the **subprocess.Popen()** function caused the upstream version of the **Eventlet** library to terminate unexpectedly. This bug has been fixed and **Eventlet** no longer fails in the described case.

BZ#960168

When a connection incoming to a server with an enabled **SSLSocket** class failed to pass the automatic **do_handshake()** function, the connection remained open. This problem affected only Python 2 versions. The underlying source code has been fixed and the failed incoming connection is now closed properly.

BZ#962779

In cases when multiple **libexpat.so** libraries were available, Python failed to choose the correct one. This update adds an explicit **RPATH** to the **_elementtree.so**, thus fixing this bug.

BZ#978129

Previously, the **urlparse** module did not parse the query and fragment parts of URLs properly for arbitrary XML schemes. With this update, **urlparse** has been fixed and correct parsing is now assured in this scenario.

Enhancement**BZ#929258**

This update adds the **collections.OrderedDict** data structure to the collections package. **collections.OrderedDict** is used in application code to ensure that the in-memory python dictionaries are emitted in the same order when converted to a string by the **json.dumps** routines.

All python users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add this enhancement.

8.161. PYTHON-BEAKER

8.161.1. RHBA-2013:1724 — python-beaker bug fix update

Updated python-beaker packages that fix one bug and are now available for Red Hat Enterprise Linux 6.

The python-beaker package provides Beaker, a web session and general caching library that includes Web Server Gateway Interface (WSGI) middleware for use in web applications.

Bug Fix

BZ#983292

Previously, Beaker used the MD5 algorithm to produce a unique hex-encoded session identifier. However, this algorithm is not, by default, supported by Python's runtime in FIPS mode. Consequently, web applications that used Beaker failed to create new sessions in certain environments. With this update, MD5 has been replaced with the SHA1 algorithm and the subsequent hex encoding with the Base64 encoding scheme. As a result, Beaker works correctly in FIPS environment.

Users of python-beaker are advised to upgrade to these updated packages, which fix this bug.

8.162. PYTHON-ETHTOOL

8.162.1. RHBA-2013:1662 — python-ethtool bug fix update

Updated python-ethtool packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The python-ethtool packages make the ethtool kernel interface available within the Python programming environment to allow querying and changing of Ethernet card settings, such as speed, port, auto-negotiation, and PCI locations.

Bug Fixes

BZ#855920

Previously, the `_haveNetwork()` routine from the `firstboot.loader` module returned "False" when set up exclusively for IPv6. As a consequence, there was no network connection. A patch that detects IPv6-only active devices has been applied, and the `_haveNetwork()` routine now returns "True" as expected.

BZ#876211

Prior to this update, if more IPv4 addresses were bound to one interface, the `pifconfig` script, which displays information about a network interface, failed to produce the correct output. As a result, two interfaces with the same IPv4 address were reported. A patch has been applied, and the `pifconfig` script now reports multiple IPv4 addresses correctly.

Users of python-ethtool are advised to upgrade to these updated packages, which fix these bugs.

8.163. PYTHON-URLGRABBER

8.163.1. RHBA-2013:1117 — python-urlgrabber bug fix update

Updated python-urlgrabber packages that fix one bug are now available.

The `python-urlgrabber` package provides `urlgrabber`, a high-level url-fetching package for the Python programming language and a corresponding utility of the same name. The `urlgrabber` package allows Python scripts to fetch data using the HTTP and FTP protocols, as well as from a local file system.

Bug Fix

BZ#807030

Previously, a flaw in the source code resulted in a traceback error when users used the `reposync` command to synchronize a remote Yum repository to a local directory, when the `utime()` system call had an error. This update corrects the mistake in the source code, and traceback errors no longer occur in the described scenario.

Users of `python-urlgrabber` are advised to upgrade to these updated packages, which fix this bug.

8.164. PYTHON-URWID

8.164.1. RHBA-2013:1550 — [python-urwid bug fix and enhancement update](#)

Updated `python-urwid` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `python-urwid` package provides a library for development of text user interface applications in the Python programming environment.



NOTE

The `python-urwid` packages have been upgraded to upstream version 1.1.1, which provides a number of bug fixes and enhancements over the previous version. Among other changes, this update resolves a number of incompatibilities with the previous version of `python-urwid` used in Red Hat Enterprise Linux 6. These incompatibilities posed a problem for Red Hat Enterprise Virtualization Hypervisor that requires the `python-urwid` packages for its new user interface. (BZ#[970981](#))

Users of `python-urwid` are advised to upgrade to these updated packages, which add these enhancements.

8.165. PYTHON-VIRTINST

8.165.1. RHBA-2013:1604 — [python-virtinst bug fix and enhancement update](#)

Updated `python-virtinst` packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The `python-virtinst` package contains several command-line utilities, including `virt-install` for building and installing new virtual machines, and `virt-clone` for cloning existing virtual machines.

Bug Fixes

BZ#861972

If there were duplicate USB devices connected to a host device, passing the vendorId or productId value to the "--host-device" option caused the virt-install utility to terminate with the following message:

```
ERROR 'vendor' and 'product', or 'bus' and 'device' are required.
```

This message did not inform about the duplicate devices, which was the cause of the error. The message has been modified to:

```
ERROR 15e1:2007 corresponds to multiple node devices
```

As a result, user is now informed about the true cause of the failed installation.

BZ#916875

Previously, when the "--disk" parameter contained the hash("#") character, incorrect error messages were displayed. This bug has been fixed and error messages now inform about the unsupported character correctly.

BZ#921480

Prior to this update, the virt-install utility created sparse logical volumes by default. However, sparse logical volumes require further configuration and therefore should be created only by advanced users or management applications. With this update, non-sparse mode is the default and only available setting.

BZ#946972

Under rare circumstances, the virt-clone utility displayed incorrect speed statistics during the clone operation. This bug has been fixed, and correct speed values are now displayed when virt-clone is in action.

BZ#954262

Under certain circumstances, when there was the ".treeinfo" file available on the system, the virt-install utility attempted to get image information from this file. In case ".treeinfo" did not contain this information, virt-install terminated unexpectedly. This bug has been fixed, and virt-install now uses default values in case of incomplete ".treeinfo".

BZ#980334

Previously, an attempt to change the model type of a video device from "qxl" led the virt-manager utility to return the following error message:

```
Error changing VM configuration: XML error: ram attribute only supported for type of qxl
```

The underlying source code has been modified, and the model type of a video device can now be changed without complications.

Enhancement

BZ#958496

This update provides a revised list of installation environments that are supported by the virt-manager utility.

Users of python-virtinst are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.166. PYTHON-WEBERROR

8.166.1. RHBA-2013:1723 — python-weberror bug fix update

An updated python-weberror package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The python-weberror package provides WebError, a web application's error handling library for use as a Web Server Gateway Interface (WSGI) middleware.

Bug Fix

BZ#746118

Previously, the WebError middleware used the MD5 algorithm when assigning an identifier to the handled error. However, this algorithm is not, by default, supported by Python's runtime in FIPS mode. Consequently, when web applications raised an exception in FIPS mode and the exception was handled by WebError, incomplete error diagnostics were provided. With this update, error identification based on MD5 is not generated automatically, thus avoiding the problems when the error identifier is not processed further.

Users of python-weberror are advised to upgrade to this updated package, which fixes this bug.

8.167. QEMU-KVM

8.167.1. RHSA-2013:1553 — Important: qemu-kvm security, bug fix, and enhancement update

Updated qemu-kvm packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having important security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

KVM (Kernel-based Virtual Machine) is a full virtualization solution for Linux on AMD64 and Intel 64 systems that is built into the standard Red Hat Enterprise Linux kernel. The qemu-kvm packages form the user-space component for running virtual machines using KVM.

CVE-2013-4344

A buffer overflow flaw was found in the way QEMU processed the SCSI "REPORT LUNS" command when more than 256 LUNs were specified for a single SCSI target. A privileged guest user could use this flaw to corrupt QEMU process memory on the host, which could potentially result in arbitrary code execution on the host with the privileges of the QEMU process.

This issue was discovered by Asias He of Red Hat.

Bug Fixes

BZ#974617

Previously, a counter variable was not correctly reset when restarting an allocating request for disk images using the *qcow2* file format. Consequently, these disk images in the cluster allocation code were corrupted in some cases. This update changes the way the number of available clusters is counted in the *qcow2* format, and *qcow2* disks are no longer corrupted in the described scenario.

BZ#927336

Due to an integer overflow in calculations, the **qemu-kvm** utility was reporting incorrect memory size on QMP (QEMU Machine Protocol) event when using **Virtio Balloon Driver** with more than 4 GB of memory. A patch has been provided to fix this bug, and **qemu-kvm** now reports the correct amount of current RAM.

BZ#917860

Previously, smart card emulation for Microsoft Windows XP and Microsoft Windows 7 guests failed due to inconsistent Answer To Reset (ATR) file length with a smart card Input/Output device error. This update creates an ATR file length with appropriate historical bytes, and disables USB signaling when necessary. Now, smart card emulation works, and failures no longer occur in the aforementioned scenario.

BZ#916020

Previously, the **qemu-kvm** utility did not enable the **IOeventFD** feature, which caused the **IOeventFD** support for **virtio-blk** devices to be silently disabled. This update enables the **IOeventFD** feature, and the **IOeventFD** support for **virtio-blk** devices works as expected.

Enhancements**BZ#670162**

A new feature for removing the backing file using the **qemu-img rebase** command has been implemented. Now, no data loss will occur when running the **qemu-img rebase** command.

BZ#963420

Red Hat Enterprise Linux 6.5 brings read-only support for VHDX (**Hyper-V** virtual hard disk), image formats, as created by Microsoft **Hyper-V**.

BZ#960685

Red Hat Enterprise Linux 6.5 brings a number of improvements on read-only support for VMDK (Virtual Machine Disk), image file formats, including its sub-formats, as created by many VMware Virtualization products.

BZ#848070

Updated support for **GlusterFS** in **QEMU** allows native access to **GlusterFS** volumes using the **libgfapi** library instead of through a locally mounted **FUSE** file system. This native approach offers considerable performance improvements.

BZ#884253

Support of Volume Control from within Microsoft Windows Guests has been implemented. Users can now fully control the volume level on Microsoft Windows XP guests using the **AC'97** codec.


BZ#914802

Support for dumping metadata of virtual disks has been implemented with this update. Third-party

applications running on the host are now able to read guest image contents without knowing the details of the *QCOW2* image format. This can be used together with the Linux device mapper to access *QCOW2* images as Linux block devices.

BZ#911569

Similarly to the Windows VSS (Visual SourceSafe) version, application-consistent snapshots can now be created with the use of scripts that attach to the **QEMU** guest agent running on the guest. These scripts can notify applications which would flush their data to the disk during a freeze or thaw operation, thus allowing consistent snapshots to be taken.

**NOTE**

VNC password authentication is disabled when the system is operating in FIPS (Federal Information Processing Standards) mode.

All `qemu-kvm` users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements. After installing this update, shut down all running virtual machines. Once all virtual machines have shut down, start them again for this update to take effect.

8.167.2. RHBA-2013:1750 — qemu-kvm bug fix update

Updated `qemu-kvm` packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

KVM (Kernel-based Virtual Machine) is a full virtualization solution for Linux on AMD64 and Intel 64 systems. `qemu-kvm` is the user-space component for running virtual machines using KVM.

Bug Fixes**BZ#1025596**

Recent changes to the block layer resulted in a disk I/O performance degradation due to the way block length is calculated and cached internally. This update improves the logic for calculating such lengths and restores performance to the expected levels.

BZ#1029327

Due to a regression, the `"qemu-img info"` command took too much time to respond with the `"cluster_size=512,preallocation=metadata"` option. This bug has been fixed and `"qemu-img info"` now responds within one second.

BZ#1029327

On images created with very small non-standard cluster sizes (for example, 512 bytes), the `"qemu-img info"` command could take a long time to respond if run immediately after an image creation. This bug has been fixed, and `"qemu-img info"` now works as expected.

BZ#1029329

When doing live migration with the `"--copy-storage-all"` option, the `virsh` user interface failed with the following error message:

```
"error: Unable to read from monitor: Connection reset by peer"
```

This bug, caused by a regression, has been fixed, and live migration now finishes successfully.

BZ#1028252

Previously, `qemu` (for example, the "`qemu-img info`" command) could not open VMWare ESX image files. A patch fixing this bug has been provided, and ESX images are now handled correctly.

Users of `qemu-kvm` are advised to upgrade to these updated packages, which fix these bugs.

8.168. QL2400-FIRMWARE

8.168.1. RHBA-2013:1707 — ql2400-firmware bug fix and enhancement update

An updated `ql2400-firmware` package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The `ql2400-firmware` package provides the firmware required to run the QLogic 2400 Series of mass storage adapters.

**NOTE**

The `ql2400-firmware` package has been upgraded to upstream version 7.00.01, which provides a number of bug fixes and enhancements over the previous version. (BZ#[996752](#))

All users of QLogic 2400 Series Fibre Channel adapters are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

8.169. QL2500-FIRMWARE

8.169.1. RHBA-2013:1709 — ql2500-firmware bug fix and enhancement update

Updated `ql2500-firmware` package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The `ql2500-firmware` package provides the firmware required to run the QLogic 2500 Series of mass storage adapters.

**NOTE**

The `ql2500-firmware` package has been upgraded to upstream version 7.00.01, which provides a number of bug fixes and enhancements over the previous version. (BZ#[996754](#))

All users of QLogic 2500 Series Fibre Channel adapters are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

8.170. QUOTA

8.170.1. RHBA-2013:1548 — quota bug fix and enhancement update

Updated quota packages that fix one bug and add two enhancements are now available for Red Hat Enterprise Linux 6.

The quota packages contain a suite of system administration tools for monitoring and limiting user and group disk usage on file systems.

Bug Fix

BZ#717948

When SELinux denied access to the quotacheck utility, the "quotacheck -c" command did not report any errors and failed with the exit code 0. With this update, quotacheck internals have been changed to propagate any error with an appropriate non-zero exit code, and to print accurate warnings if the old quota file could not be used. As a result, quotacheck now reports errors correctly while initializing quotas on a file system.

Enhancements

BZ#890051

The edquota and setquota utilities now support input values entered with the use of multiplicative unit abbreviations (K, M, G, T).

BZ#846120

With this update, the quotacheck utility no longer checks quota files on the GFS2 and XFS file systems. Since these file systems do not have any quota files, the check was redundant.

Users of quota are advised to upgrade to these updated packages, which fix this bug and add these enhancements.

8.171. RDESKTOP

8.171.1. RHBA-2013:1556 — rdesktop bug fix and enhancement update

Updated rdesktop packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The rdesktop package is an open source client for Windows NT Terminal Server and Windows 2000 & 2003 Terminal Services, capable of natively speaking Remote Desktop Protocol (RDP) in order to present the user's NT desktop.



NOTE

The rdesktop package has been upgraded to upstream version 1.7.0., which provides a number of bug fixes and enhancements over the previous version, including fixes for incorrect cursor and scrollbar rendering. (BZ#852890, BZ#914279)

Bug Fix

BZ#902912

Due to a bug in the rdesktop code, error messages were continuously displayed in the shell while rdesktop was used. This bug has been fixed and messages are no longer generated during an rdesktop session.

Enhancement

BZ#701246

With this update, support for Terminal Services on clustered Windows Server 2008 R2 machines has been added to the rdesktop package.

Users of rdesktop are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.172. RDMA STACK

8.172.1. RHSA-2013:1661 — Moderate: RDMA stack security, bug fix, and enhancement update

Updated rdma, libibverbs, libmlx4, librdmacm, qperf, perftest, openmpi, compat-openmpi, infinipath-psm, mpitests, and rds-tools packages that fix two security issues, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE link(s) associated with each description below.

Red Hat Enterprise Linux includes a collection of Infiniband and iWARP utilities, libraries and development packages for writing applications that use Remote Direct Memory Access (RDMA) technology.

Security Fixes

CVE-2013-2561

A flaw was found in the way ibutils handled temporary files. A local attacker could use this flaw to cause arbitrary files to be overwritten as the root user via a symbolic link attack.

CVE-2012-4516

It was discovered that librdmacm used a static port to connect to the ib_acm service. A local attacker able to run a specially crafted ib_acm service on that port could use this flaw to provide incorrect address resolution information to librdmacm applications.

The CVE-2012-4516 issue was discovered by Florian Weimer of the Red Hat Product Security Team.

This advisory updates the following packages to the latest upstream releases, providing a number of bug fixes and enhancements over the previous versions:

Table 8.1. Upgraded packages

Package name	Upstream version
libibverbs	1.1.7
libmlx4	1.0.5

Package name	Upstream version
librdmacm	1.0.17
mstflint	3.0
perftest	2.0
qperf	0.4.9
rdma	3.10

Several bugs have been fixed in the `openmpi`, `mpitests`, `ibutils`, and `infinipath-psm` packages.

The most notable changes in these updated packages from the RDMA stack are the following:

- Multiple bugs in the Message Passing Interface (MPI) test packages were resolved, allowing more of the `mpitest` applications to pass on the underlying MPI implementations.
- The `libmlx4` package now includes `dracut` module files to ensure that any necessary custom configuration of `mlx4` port types is included in the `initramfs` `dracut` builds.
- Multiple test programs in the `perftest` and `qperf` packages now work properly over RoCE interfaces, or when specifying the use of `rdmacm` queue pairs.
- The `mstflint` package has been updated to the latest upstream version, which is now capable of burning firmware on newly released Mellanox Connect-IB hardware.
- A compatibility problem between the `openmpi` and `infinipath-psm` packages has been resolved with new builds of these packages.

All RDMA users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.173. READAHEAD

8.173.1. [RHBA-2013:1739](#) — [readahead bug fix update](#)

Updated `readahead` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `readahead` packages provide an utility that reads the contents of a list of files into memory. The files are read from the cache directory when they are actually needed. Its goal is to speed up the boot process.

Bug Fix

BZ#[1017072](#)

The `readahead-collect` tool of the `readahead` package delayed the startup process of the `auditd` daemon on system boot. This could have caused various problems with auditing, `kdump` and `runlevel` switching. To speed up the boot process, `readhead` is now turned off by default. If needed, `readhead` can be turned on by editing the `/etc/sysconfig/readahead` file.

Users of readahead are advised to upgrade to these updated packages, which fix this bug.

8.174. REDHAT-INDEXHTML

8.174.1. RHEA-2013:1725 — [redhat-indexhtml enhancement update](#)

Updated redhat-indexhtml package that adds two enhancements is now available for Red Hat Enterprise Linux 6.

The redhat-indexhtml package contains a welcome page which is displayed in the web browser after successful installation of Red Hat Enterprise Linux. This web page provides information on registration, documentation, and support for Red Hat Enterprise Linux.

Enhancements

BZ#892016

With this update, links in the index.html file have been updated to point directly to the Red Hat Customer Portal and Red Hat Enterprise Linux 6 Release Notes. Also, several links on translated versions of index.html were modified to point to Release Notes in the corresponding language.

BZ#657558

Translations of index.html have been added for Assamese, Bengali, and Telugu languages.

Users of redhat-indexhtml are advised to upgrade to these updated packages, which add these enhancements.

8.175. REDHAT-RELEASE

8.175.1. RHEA-2013:1546 — [redhat-release enhancement update for Red Hat Enterprise Linux 6.5](#)

Enhanced redhat-release packages are now available for Red Hat Enterprise Linux 6.5.

The redhat-release package contains licensing information regarding, and identifies the installed version of, Red Hat Enterprise Linux.

These updated redhat-release packages reflect changes made for the release of Red Hat Enterprise Linux 6.5.

Users of Red Hat Enterprise Linux 6 are advised to upgrade to these updated redhat-release packages, which add this enhancement.

8.176. RED HAT ENTERPRISE LINUX 6.5 RELEASE NOTES

8.176.1. RHEA-2013:1717 — [Red Hat Enterprise Linux 6.5 Release Notes](#)

Updated packages containing the Release Notes for Red Hat Enterprise Linux 6.5 are now available.

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security and bug fix errata. The Red Hat Enterprise Linux 6.5 Release Notes documents the major changes made to the Red Hat Enterprise Linux 6 operating system and its accompanying applications for this minor

release. Detailed notes on all changes in this minor release are available in the Technical Notes.

Refer to the Online Release Notes for the most up-to-date version of the Red Hat Enterprise Linux 6.5 Release Notes:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html-single/6.5_Release_Notes/index.html

8.177. RESOURCE-AGENTS

8.177.1. RHBA-2013:1541 — resource-agents bug fix and enhancement update

Updated resource-agents packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The resource-agents packages contain a set of scripts to interface with several services to operate in a High Availability environment for both Pacemaker and rgmanager service managers.

Bug Fixes

BZ#784933

Previously, when the **exportfs** utility was used to relocate an exported share, the size of the **/var/llib/nfs/rmtab** file was doubled. This bug has been fixed and the **/var/lib/nfs/rmtab** file size is no longer doubled in the aforementioned scenario.

BZ#851188

Prior to this update, the **fs-lib.sh** agent did not recognize the trailing slash ("/") character when searching for devices in the **/proc/mounts** file. Consequently, NFSv4 mounts were not monitored. With this update, **fs-lib.sh** has been modified to track the slash characters. As a result, NFSv4 mounts are managed and monitored as expected.

BZ#853220

Due to a bug in the **oracledb.sh** script, when there were multiple ORACLE instances running in the same home directory, the script produced unnecessary delays. The bug has been fixed, and **oracledb.sh** now works without delays when multiple ORACLE instances are present in the home directory.

BZ#871659

To shutdown cleanly, the **postgres** agent needs to receive the SIGINT signal. Previously, this signal was not sent and **postgres** performed a hard shutdown instead of a graceful exit. This behavior has been modified, and SIGINT is now sent to **postgres** on shutdown to attempt a graceful exit, and after a period of time, the SIGQUIT signal is sent if the agent is still active. As a result, **postgres** performs graceful shutdown during the stop action.

BZ#884326

Previously, if a device failed in a non-redundant (that is not mirror or RAID) logical volume (LV) that was controlled by the **HA-LVM** utility, the entire LV could be automatically deleted from the volume group. This bug has been fixed and now, if a non-redundant logical volume suffers a device failure, **HA-LVM** fails to start the service rather than forcing the removal of failed LVs from the volume group.

BZ#895075

Prior to this update, the **ip.sh** agent did not configure IPv6 addresses that contained upper-case letters. Consequently, a resource with such an address failed. With this update, **ip.sh** has been modified to be case insensitive for IPv6 addresses. As a result, IPv6 addresses with upper case letters are now configured properly by **ip.sh**.

BZ#908457

Previously, agents based on the **fs-lib.sh** script, such as **ip.sh**, ignored the **self_fence** option when the **force_unmount** option was enabled. Consequently, the configured **self_fence** option was not enabled. This bug has been fixed and **self_fence** is accepted regardless of **force_unmount**.

BZ#948730

With this update, the priority level of log messages produced by the **mount** utility has been changed from previous **error** to more appropriate **debug** level.

BZ#959520

Due to an incorrect SELinux context of the **/var/lib/nfs/statd/sm/** directory, the **rpc.statd** daemon was unable to start. This problem only appeared if the cluster included NFS mounts. This update modifies how files are copied to the **/var/lib/nfs/statd/sm/** directory, so that the SELinux context is inherited from the target directory. As a result, **rpc.statd** can now be started without complications.

BZ#974941

When **autofs** maps are used for network storage, agents for cluster file systems ("fs") such as **netfs.sh**, **fs.sh**, or **clusterfs.sh** require the **use_findmnt** option set to **'false'**. Previously, when **use_findmnt** was set incorrectly, and **autofs** maps became unavailable, the **rgmanager** services with "fs" resources consequently became unresponsive until the network was restored. The underlying source code has been modified and **rgmanager** services no longer hang in the aforementioned scenario.

BZ#976443

Prior to this update, the **lvm.sh** agent was unable to accurately detect a tag represented by a cluster node. Consequently, the active logical volume on a cluster node failed when another node rejoined the cluster. With this update, **lvm.sh** properly detects whether tags represent a cluster node. As a result, when nodes rejoin the cluster, the volume group no longer fails on other nodes.

BZ#981717

When multiple instances of the **tomcat-6** service were used as cluster resources, the **TOMCAT_USER** setting in custom **/conf/tomcat6.conf** configuration files was ignored. Consequently, each instance always started with **TOMCAT_USER** set to **root**. This bug has been fixed, and **TOMCAT_USER** is now applied properly in the described case.

BZ#983273

Under certain circumstances, when the **tomcat.conf** configuration file for a **tomcat-6** resource was stored on a shared storage resource that became unavailable, the subsequent stop operation on **tomcat-6** failed. This bug has been fixed, and **tomcat-6** can now be successfully stopped when **tomcat.conf** is not readable.

BZ#998012

File system based resources, such as **fs.sh** or **clusterfs.sh**, required usage of the **/tmp** directory during status monitoring. If this directory became full after mounting the file system, the monitor action failed even though the file system was correctly mounted. The **/tmp** directory is no longer used during file system monitors, thus fixing this bug.

BZ#1009772

If **rgmanager** was started simultaneously on two nodes, these nodes could both execute the **lvchange --deltag** command at the same time and corrupt the LVM headers. With this update, LVM headers do not become corrupt even when **rgmanager** starts on two nodes at the same time.

BZ#1014298

Previously, when the NFS server was unresponsive, the **fuser** utility could block the unmounting of an NFS file system. With this update, **fuser** has been replaced with custom logic that searches for processes with open file descriptors to an NFS mount, thus fixing this bug.

Enhancements**BZ#670022**

With this update, support for Oracle Database 11g has been added to the **oracledb**, **orainstance**, and **oralistener** resource agents.

BZ#711586

This update adds the new **update-source** option to the **named.sa** agent. With this option enabled, it is possible to set the **notify-source**, **transfer-source**, and **query-source** to the service cluster IP.

BZ#909954

With this update, the lock file for the **/usr/share/cluster/orainstance.sh** script has been moved from the **/tmp/** directory to **/var/tmp/**.

BZ#917807

With this update, the **TNS_ADMIN** variable has been added to the **oracledb.sh** cluster script. This variable is a standard Oracle feature to set a specific path to the listener configuration file.

BZ#919231

This update improves performance of start, stop, and monitor actions for file system resources. The file system resources use of the **findmnt** utility to speed up migrations on clusters with a large amount of file system resources.

BZ#989284

This update adds official support for the ocf heartbeat resource agents required for use with the Pacemaker cluster manager. Only officially supported agents are present for this initial release. This means heartbeat agents lacking official support are not shipped in this update.

Users of resource-agents are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.177.2. RHBA-2013:1746 — resource-agents bug fix update

Updated resource-agents packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The resource-agents packages contain a set of scripts to interface with several services to operate in a High Availability environment for both Pacemaker and rgmanager service managers.

Bug Fixes

BZ#1027410

Prior to this update, the netfs agent could hang during a stop operation, even with the self_fence option enabled. With this update, self fence operation is executed sooner in the process, which ensures that NFS client detects server leaving if umount can not succeed, and self fencing occurs.

BZ#1027412

Previously, the IPaddr2 agent did not send out unsolicited neighbor advertisements to announce a link-layer address change. Consequently, floating IPv6 addresses, which require this functionality, could not work correctly. To fix this bug, the send_ua internal binary required for IPaddr2 agent to drive IPv6 addresses has been added. As a result, the floating IPv6 addresses now work correctly, and IPv4 addresses are left unaffected by this change.

Users of resource-agents are advised to upgrade to these updated packages, which fix these bugs.

8.178. RGMANAGER

8.178.1. RHBA-2013:1600 — rgmanager bug fix update

Updated rgmanager packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The rgmanager package contains the Red Hat Resource Group Manager, which is used to create and manage high-availability server applications in the event of system downtime.

Bug Fixes

BZ#862075

Previously, if the main rgmanager process died, either by an unexpected termination with a segmentation fault or by killing it manually, any service running on it, was immediately recovered on another node rather than waiting for fencing, like the rgmanager process did in previous versions. This was problematic for services containing Highly Available Logical Volume Manager (HA-LVM) resources using tagging, because the start operation failed if the tag that was found belonged to a node that was still a member of the cluster. With this update, service recovery is delayed until after the node is removed from the configuration and fenced, which allows the LVM resource to recover properly.

BZ#983296

Previously, attempts to start an MRG Messaging (MRG-M) broker caused rgmanager to terminate unexpectedly with a segmentation fault. This was caused by subtle memory corruption introduced by calling the pthread_mutex_unlock() function on a mutual exclusion that was not locked. This update addresses scenarios where memory could be corrupted when calling pthread_mutex_unlock(), and rgmanager no longer terminates unexpectedly in the described situation.

Users of rgmanager are advised to upgrade to these updated packages, which fix these bugs.

8.179. RHEL-GUEST-IMAGE

8.179.1. RHBA-2013:1735 — rhel-guest-image bug fix and enhancement update

Updated rhel-guest-image packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The rhel-guest-image packages provide a Red Hat Enterprise Linux 6.5 KVM Guest Image for cloud instances. This image is provided as a minimally configured system image which is available for use as-is or for configuration and customization as required by end users.

Bug Fixes

BZ#912475, BZ#952280

Prior to this update, the `/etc/ssh/sshd_config` file was not created properly due to a missing End of Line (EOL) return. Consequently, the `sshd` daemon start failed with an error message. The writing of the sequence to `/etc/ssh/sshd_config` has been corrected, and `sshd` start no longer fails in this scenario.

BZ#912801

Previously, the persistent `udev` rule was not removed from the guest image when the virtual machine (VM) was shut down. Consequently, when the VM was booted again and a different network card media access control (MAC) address was assigned to it, the `udev` rule caused the Network Interface Controller (NIC) to be set as `eth1` instead of `eth0`. This incorrect configuration caused the instance to boot with no network support. This update adds the `/etc/udev/rules.d/75-persistent-net-generator.rules` file, and the VM configuration works as expected in the described scenario.

BZ#969487

The Red Hat Enterprise Linux guest image did not make use of the `ttyS0` serial port by default, so tools that monitor the serial console log did not capture any information. This update adds a console log feature to the boot loader, and the Red Hat Enterprise Linux guest image now prints boot messages on the `ttyS0` serial port and the standard console.

BZ#983611

Metadata access is handled by the network node, and the `cloud-init` service uses the metadata at startup time. Previously, the `cloud-init` service did not have the `"NOZEROCONF=yes"` stanza configured. Consequently, access to the subnet `169.254.0.0/16` range was not routed to the network node, so the `cloud-init` service failed to work. This update adds `"NOZEROCONF=yes"` to cloud images in the `/etc/sysconfig/network` file. Users should avoid turning on the zeroconf route in these images. To disable the zeroconf route at system boot time, edit the `/etc/sysconfig/network` file as root and add `"NOZEROCONF=yes"` to a new line at the end of the file.

BZ#1006883

Previously, the `/etc/fstab` file contained unnecessary entries. Consequently, file systems mounting did not work properly when updating the kernel and grub installation. With this update, the unnecessary entries have been removed and file systems mounting now works as expected in this scenario.

BZ#1011013

Previously, a `dhclient` configuration option was missing in the `ifcfg-eth0` configuration file. As a consequence, the network connectivity was lost if Dynamic Host Configuration Protocol (DHCP) failed. This update adds the `"PERSISTENT_DHCLIENT=1"` configuration option to the `ifcfg-eth0`

configuration file. Now, if the Network Interface Controller (NIC) is unable to negotiate a DHCP address, it retries.

Enhancement

BZ#974554

The virtual-guest image now includes a profile for the "tuned" daemon that is activated by default. This helps improve performance of the guest in most situations.

Users of rhel-guest-image are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.180. RHN-CLIENT-TOOLS

8.180.1. RHBA-2013:1702 — rhn-client-tools bug fix update

Updated rhn-client-tools packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

Red Hat Network Client Tools provide programs and libraries that allow the system to receive software updates from Red Hat Network (RHN).

Bug Fixes

BZ#891746

Previously, the rhn-channel manual page incorrectly referred to the "--username" option instead of "--user". This mistake has been corrected and the rhn-channel manual page now correctly refers to the "--user" option.

BZ#912984

Prior to this update, some messages written in English occurred in the Japanese installation of Red Hat Enterprise Linux 6.4. The untranslated strings have been translated, and the messages are now shown in the correct language.

BZ#983999

Previously, the rhn-client-tools code called by the sosreport utility on Red Hat Enterprise Linux 6 terminated unexpectedly with a traceback. The bug has been fixed and the information about hardware is now correctly gathered by sosreport.

BZ#994531

Previously, a machine with many CPUs could report large value for idle time for all its processors. Consequently, the idle time value did not fit into XML-RPC's integer limits and running the rhn_check command on a problematic machine resulted in a traceback error. The bug has been fixed and rhn_check in the problematic scenarios works now correctly.

BZ#997637

Previously, the rhn-profile-sync utility terminated unexpectedly with a traceback when an older version of the rhn-virtualization-host package was installed on the machine. The bug has been fixed by requiring a newer version of rhn-virtualization-host.

Users of `rhn-client-tools` are advised to upgrade to these updated packages, which fix these bugs.

8.180.2. [RHBA-2013:1087](#) — `rhn-client-tools` bug fix and enhancement update

Updated `rhn-client-tools` packages that fix one bug and add one enhancement are now available.

Red Hat Network Client Tools provide programs and libraries that allow systems to receive software updates from Red Hat Network (RHN).

Bug Fix

BZ#949648

The RHN Proxy did not work properly if separated from a parent by a slow enough network. Consequently, users who attempted to download larger repodata files and RPMs experienced timeouts. This update changes both RHN Proxy and Red Hat Enterprise Linux RHN Client to allow all communications to obey a configured timeout value for connections.

Enhancement

BZ#949640

While Satellite 5.3.0 now has the ability to get the number of CPUs via an API call, there was no function to obtain the number of sockets from the registered systems. This update adds a function to get the number of physical CPU sockets in a managed system from Satellite via an API call.

Users of `rhn-client-tools` are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

8.181. RHNLIB

8.181.1. [RHBA-2013:1085](#) — `rhnlib` bug fix and enhancement update

Updated `rhnlib` packages that fix one bug are now available.

The `rhnlib` packages contain Python libraries developed specifically for interfacing with the Red Hat Network.

Bug Fix

BZ#949650

The RHN Proxy did not work properly if separated from a parent by a slow enough network. Consequently, users who attempted to download larger repodata files and RPMs experienced timeouts. This update changes both RHN Proxy and Red Hat Enterprise Linux RHN Client to allow all communications to obey a configured timeout value for connections.

Users of `rhnlib` are advised to upgrade to these updated packages, which fix this bug.

8.182. RICCI

8.182.1. [RHBA-2013:1673](#) — `ricci` bug fix and enhancement update

Updated ricci packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The ricci packages contain a daemon and a client for remote configuring and managing of clusters.

Bug Fixes

BZ#853890

Prior to this update, the `ccs_sync` command, used to copy and synchronize the cluster configuration across different nodes, could terminate unexpectedly when pushing an external configuration file to selected nodes. With this update, `ccs_syncd` has been fixed and it no longer crashes in the aforementioned scenario.

BZ#883585

Prior to this update, the `ricci`, `ccs`, and `ccs_sync` man pages had file permissions incorrectly set to executable. With this update, the file permissions of these man pages have been corrected.

BZ#893574

Prior to this update, the `ccs` manager did not configure the `fence_scsi` agent properly. The missing configuration could cause an incorrect behavior of a cluster during recovery. This bug has been fixed and `fence_scsi` is now configured as expected.

BZ#918555

Due to an incorrect parsing of cluster configuration, the `ccs` manager did not display the `fence_daemon` options. With this update, the parsing has been fixed and `fence_daemon` options are now properly displayed.

Enhancements

BZ#877863

With this update the `--lsresourceopts` command-line option has been added to the `css` manager. This option provides an intuitive way how to display and list all available resource types and agents from the shell prompt.

BZ#1009098

The cluster schema included in the `ricci` package has been updated to include new options in resource and fence agents packages.

Users of `ricci` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.183. RP-PPPOE

8.183.1. RHBA-2013:0952 — rp-pppoe bug fix update

Updated `rp-pppoe` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `rp-pppoe` packages provide the Roaring Penguin PPPoE (Point-to-Point Protocol over Ethernet) client, a user-mode program that does not require any kernel modifications. This client is fully compliant with RFC 2516, the official PPPoE specification.

Bug Fix

BZ#841190

Previously, the `pppoe-server` service started by default at each system boot, which was not intended as `pppoe-server` is supposed to run only when enabled by an administrator. This update ensures that `pppoe-server` is not started by default, thus fixing this bug.

Users of `rp-pppoe` are advised to upgrade to these updated packages, which fix this bug.

8.184. RPM

8.184.1. RHBA-2013:1665 — rpm bug fix update

Updated rpm packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The RPM Package Manager (RPM) is a command-line driven package management system capable of installing, uninstalling, verifying, querying, and updating software packages.

Bug Fixes

BZ#868332

Previously, the `brp-python-bytecompile` script skipped those paths that included `"/usr/lib.*/python.+/"` string. Consequently, when creating an RPM that contained Python modules in paths (like `"/opt/myapp/usr/lib64/python2.6/site-packages/mypackage/"`), bytecode was not created. The reproducer specification has been changed, and bytecode is now created for all paths.

BZ#904818

When wildcard characters were used with the `"%caps"` tag in the spec file, the `rpmbuild` utility terminated unexpectedly. The provided patch corrects the problem by making a copy of the caps data for each file it applies to, and thus `rpmbuild` no longer crashes in the described scenario.

BZ#919435

Previously, when installing a package with a high (80k) number of files, the RPM Package Manager terminated unexpectedly with a segmentation fault. As a workaround, the segmentation fault has been replaced with an error message when a package with a high number of files fails to install.

BZ#920190

Previously the `rpm` program attempted to unconditionally process any `"%include"` directive it found in a spec file, either leading to unwanted content in the package or error messages. The updated rpm package properly honours various spec conditionals for `"%include"`.

BZ#963724

With this update, Red Hat Enterprise Linux 5 backwards-compatibility option, `"%_strict_script_errors macro"`, has been added. The default behavior of Red Hat Enterprise Linux 6 does not change with this update and users that do not demand this option specifically are not advised to use it.

Users of `rpm` are advised to upgrade to these updated packages, which fix these bugs. All running applications linked against the RPM library must be restarted for this update to take effect.

8.185. RPMLINT

8.185.1. RHBA-2013:0948 — rpmlint bug fix update

Updated rpmlint packages that fix two bugs are now available.

Rpmlint is a tool for checking common errors in rpm packages. It can be used to test individual packages and specific files before uploading or to check an entire distribution. By default all applicable checks are processed but specific checks can be performed by using command line parameters.

Bug Fixes

BZ#663082

Previously, an incorrect rule in rpmlint caused it to report a "missing-lsb-keyword Default-Stop" error message for services that did not start by default in any run level. This update corrects the rule in rpmlint, and error messages no longer occur in the described scenario.

BZ#958038

When a package had a UTF-8 encoding error in the description, or the description was encoded in a different character set, then rpmlint terminated unexpectedly with a segmentation fault. With this update, rpmlint returns an error message that the description has incorrectly encoded UTF-8 data (tag-not-utf8 error), and crashes no longer occur in the described scenario.

Users of rpmlint are advised to upgrade to these updated packages, which fix these bugs.

8.186. RSYSLOG

8.186.1. RHBA-2013:1716 — rsyslog bug fix update

Updated rsyslog packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The rsyslog packages provide an enhanced, multi-threaded syslog daemon. It supports MySQL, syslog/TCP, RFC 3195, permitted sender lists, filtering on any message part, and fine grain output format control.

Bug Fixes

BZ#862517

The imgssapi module is initialized as soon as the configuration file reader encounters the `$InputGSSServerRun` directive in the `/etc/rsyslog.conf` configuration file. The supplementary options configured after `$InputGSSServerRun` are therefore ignored. For configuration to take effect, all imgssapi configuration options must be placed before `$InputGSSServerRun`. Previously, when this order was reversed, the rsyslogd daemon terminated unexpectedly with a segmentation fault. This bug has been fixed, and rsyslogd no longer crashes in the described scenario.

BZ#886117

Rsyslog directives used for controlling the file owner or group (`FileOwner`, `FileGroup`, `DirOwner`, `DirGroup`) translate names to numerical IDs only during rsyslog's initialization. Previously, when user data were not available at rsyslog's startup, IDs were not assigned to these log files. With this

update, new directives that do not depend on the translation process have been added (FileOwnerId, FileGroupId, DirOwnerId, DirGroupId). As a result, log files are assigned the correct user or group ID even when user information is not available during rsyslog's startup.

BZ#893197

Due to a bug in the source code, the host name was replaced by an empty string if the \$RepeatedMsgReduction directive was enabled. This bug has been fixed, and the host name is now stored correctly when \$RepeatedMsgReduction is on.

BZ#924754

Prior to this update, the \$FileGroup directive did not process groups larger than a certain size. Consequently, when this size was reached, the rsyslogd daemon failed to set the requested group and the root user was left as the owner of a file. This bug has been fixed and \$FileGroup now creates groups properly in the described case.

BZ#927405

An erroneous patch in a previous release, which changed the implementation of the configuration file parser, caused the rsyslogd daemon to terminate unexpectedly with a segmentation fault for certain configurations. With this update, the patch has been removed, and file crashes no longer occur with the default configuration. However, the \$IncludeConfig directive must be placed at the beginning of the /etc/rsyslog.conf configuration file before other directives. If there is need to use \$IncludeConfig further in the file, users are advised to prepend it with a dummy action such as "syslog.debug /dev/null".

BZ#951727

Prior to this update, a numerical value of the PRI property was appended to the pri-text variable. The resulting pri-text value looked for example like "local0.info≤164>". With this update the suffix has been removed. Now, the variable only contains textual facility and severity values.

BZ#963942

Previously, an incorrect data type was set for the variable holding the spool file size limit. Consequently, the intended size limit was not accepted and a message loss could occur. With this update, the data type of the aforementioned variable has been corrected. As a result, spool files are set correctly with the user-defined size limit.

Users of rsyslog are advised to upgrade to these updated packages, which fix these bugs.

8.187. RUBYGEMS

8.187.1. RHBA-2013:1694 — rubygems bug fix and enhancement update

Updated rubygems packages that fix one bug and add one enhancement are now available for Red Hat Enterprise Linux 6.

RubyGems is the Ruby standard for publishing and managing third-party libraries.

Bug Fix

BZ#559707

Previously, the specification file listed an incorrect license. The specification file has been updated to fix the license, which is MIT now.

Enhancement

BZ#788001

New release of `rubygems` package introduces `rubygems-devel` subpackage with RPM macros for easier packaging and better compatibility with Fedora.

Users of `rubygems` are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

8.188. S390UTILS

8.188.1. RHBA-2013:1678 — s390utils bug fix and enhancement update

Updated `s390utils` packages that fix a number of bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `s390utils` packages contain a set of user space utilities for Linux on IBM System z architecture.

Bug Fixes

BZ#883456

Previously, the `ziomon` utility did not follow symbolic links to find multipath devices in the `/dev/mapper/` directory. Consequently, the multipath devices could not be found. The bug has been fixed with this update so that `ziomon` now follows the symbolic links and the multipath devices can be found as expected.

BZ#887336

The `dbginfo.sh` utility collects various data from the system for debugging purposes. Previously, certain runtime data were missing from the `dbginfo.sh` output and the underlying source code was not coherent. As a consequence, incomplete information was provided and the utility performance was decreased. In addition, in certain cases, `dbginfo.sh` failed to detect if the `debugfs` file system had been mounted. The code has been unified and calls to additional utilities and commands have been added to improve collecting data. Also, `dbginfo.sh` now collects data from additional configuration and log files.

BZ#906836

The `ziorep_config` configuration report is supposed to ignore Small Computer System Interface (SCSI) disks that are not part of the multipath devices when creating the multipath mapper report. Previously, `ziorep_config` failed to correctly ignore SCSI disks, which were not a part of a multipath device. Now, when no multipath device is found for a SCSI disk, such a disk is skipped in the output.

BZ#948343

Previously, the `sysfs_getUnitsFromPort()` function only searched the Small Computer System Interface (SCSI) device directory for devices using the `scsi_generic:sg*` layout. This layout is deprecated and available only if the `CONFIG_SYSFS_DEPRECATED[_V2]` option is set in the kernel

configuration. Consequently, the function did not work properly. With this update, the function has been modified to search for devices using also the **scsi_generic/sg*** layout so that it now works as expected.

BZ#951585

The World Wide Names (WWNs) and Logic Unit Numbers (LUNs) strings, and the `fc_host` statistics `sysfs` attributes were converted to an incorrect integer type, which was too small to hold the entire possible range of the strings. This could cause a loss of information. The underlying source code has been modified to fix this bug and WWNs, LUNs, and the `fc_host` statistics `sysfs` attributes are now converted to the correct integer type.

BZ#973235

Due to the incoherent **dbginfo.sh** source code, the collection of the `sysfs` tree took a long time and logs were not written serialized, but were mixed up. Also, some information was missing from the generated file, because the utility did not collect information from all necessary configuration files. With this update, the underlying source code has been improved to fix these problems and **dbginfo.sh** now works as expected.

BZ#974180

The **dbginfo.sh** utility collects various data from the system for debugging purposes. Previously, the collected information from the system did not provide enough data about cryptographic adapters. The **dbginfo.sh** has been modified to collect information providing further information about the adapters.

BZ#996180

The trace pipes for CPU tracing in the `sysfs` file system could potentially block the **dbginfo.sh** utility. Consequently, the utility became unresponsive in such a case. This bug has been fixed so that the trace pipes no longer block **dbginfo.sh**. As a result, the utility no longer hangs in the described scenario.

BZ#997359

The **zgetdump** utility did not allocate enough memory for the CPU ELF notes. Consequently, on systems with many CPUs, the following error was returned:

```
zgetdump: Internal Error: hdr_size=28512 alloc_size=26624
```

With this update, the utility has been modified to allocate enough memory for the ELF notes and the error is no longer returned.

BZ#997360

Previously, the even addresses were loaded incorrectly. Consequently, when the **--force** option for the Direct Access Storage Device (DASD) multi-volume dump had been specified and the dump partition was modified afterwards, the dump failed with an error. With this update, the correct even address is loaded and the option works as expected.

Enhancements**BZ#929261**

This enhancement provides the Fuzzy live dump feature, which is a utility that can extract the current memory state of the kernel. To do so, Fuzzy live dump extracts an ELF core dump and filters it with

the **makedumpfile** command. This feature allows users to provide a problem analysis without shutting down the system.



NOTE

The recorded memory can change during recording so that results may not be consistent in all cases.

BZ#929263

This enhancement introduces a new Direct Access Storage Device (DASD) interface, which writes all outstanding data to a DASD device before setting it offline. Users now can use the **safe offline** option to ensure that all outstanding write requests are completed before setting the device offline.

BZ#967014

Support for the physical channel-ID (PCHID) mapping has been added to the s390utils packages allowing users to determine PCHID associated with a logical channel-path identifier (CHPID). The ability to map CHPID to PCHID is important for maintenance purposes and error determination processes. In addition, advanced health checks, which rely on the PCHID information, can be now enabled.

Users of s390utils are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.189. SAMBA

8.189.1. [RHSA-2013:1542](#) — **Moderate: samba security, bug fix, and enhancement update**

Updated samba packages that fix three security issues, several bugs, and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

Samba is an open-source implementation of the Server Message Block (SMB) or Common Internet File System (CIFS) protocol, which allows PC-compatible machines to share files, printers, and other information.

Security Fixes

[CVE-2013-0213](#)


It was discovered that the Samba Web Administration Tool (SWAT) did not protect against being opened in a web page frame. A remote attacker could possibly use this flaw to conduct a clickjacking attack against SWAT users or users with an active SWAT session.

[CVE-2013-0214](#)

A flaw was found in the Cross-Site Request Forgery (CSRF) protection mechanism implemented in SWAT. An attacker with the knowledge of a victim's password could use this flaw to bypass CSRF protections and conduct a CSRF attack against the victim SWAT user.

CVE-2013-4124

An integer overflow flaw was found in the way Samba handled an Extended Attribute (EA) list provided by a client. A malicious client could send a specially crafted EA list that triggered an overflow, causing the server to loop and reprocess the list using an excessive amount of memory.



NOTE

This issue did not affect the default configuration of the Samba server.

Red Hat would like to thank the Samba project for reporting CVE-2013-0213 and CVE-2013-0214. Upstream acknowledges Jann Horn as the original reporter of CVE-2013-0213 and CVE-2013-0214.

Bug Fixes

BZ#948071, BZ#953985

An attempt to retrieve group information from a trusted domain using a connection based on the TCP/IP protocol failed, because the machine account credentials were required to establish a secured connection over TCP/IP. Consequently, a fallback to a named pipe connection did not work, and users were not able to log into a trusted domain. With this update, the fallback to a named pipe connection has been fixed and users can now log into trusted domains as expected.

BZ#951175


Previously, when the **winbind** service (**winbindd**) was under a heavy load to authenticate a large amount of Active Directory (AD) users, it was possible that it used 100% of the CPU and stopped the user authentication. This update provides a patch to improve the connection handling significantly, and **winbindd** no longer stops the user authentication in the described scenario.

BZ#952268

The Samba service contains the user name mapping optimization that stores an unsuccessful mapping so that it is not necessary to traverse the whole mapping file every time. Due to a bug in the optimization, the user name mapping worked only once and then was subsequently overwritten by an unsuccessful mapping. This update provides a patch to fix this bug and the successful user name mapping is no longer overwritten in the described scenario.

BZ#953025

Previously, guest users in the "security = share" mode did not have the correct token that allowed write operations on a writable guest share. Consequently, such users were not able to create or write to any files within the share. With this update, a patch has been provided to fix this bug and the guest users are able to write to or create any files within the writable share as expected.



NOTE

The "security = share" mode is deprecated and users should migrate to the "security = user" mode.

BZ#955683

The **net ads keytab add** command always converted characters in the service principal name (SPN) into uppercase characters. Consequently, several Kerberos services were not able to find their tickets. With this update, SPN is no longer converted into uppercase characters and Samba works as

expected.

BZ#961932

Due to a bug in the authentication code that forwarded the NTLMv2 authentication challenge to the primary domain controller (PDC), an incorrect domain name could be sent from a client. Consequently, the user was not able to log in because when the domain name was hashed in the second NTLMv2 authentication challenge, the server could not verify the validity of the hash and the access was rejected. With this update, the correct domain name is set by the client to the PDC and the user is able to log in as expected.

BZ#980382

An attempt to execute the **wkssvc_NetWkstaEnumUsers** RPC command without a pointer to the resume handle caused the **smbd** daemon to terminate with a segmentation fault. Consequently, the client was disconnected. With this update, the underlying source code has been adapted to verify that the pointer is valid before attempting to dereference it. As a result, **smbd** no longer crashes in this situation.

BZ#997338

When a non-root user executed the **smbstatus** command, the locked files were missing from the command output. The underlying source code has been modified to fix this bug and non-root users are now able to display the locked files as expected.

BZ#1003689

Red Hat Enterprise Linux 6 can be used as a print server that shares network printers used by Microsoft Windows 8 clients. Previously, the version of Samba shipped with Red Hat Enterprise Linux was not compatible with Windows 8. Consequently, when a Windows 8 client accessed a printer share and attempted to install the driver for this printer, an error occurred. This update applies a patch to fix this bug and Windows printer drivers can be now installed successfully in the described scenario.

BZ#1008574

Previously, the main **winbind** daemon was not informed when its child process had successfully connected to a domain controller. As a consequence, the Network Data Representation (NDR) cache entries never expired and therefore the entries could not be updated. With this update, the **winbind** child process notifies the main **winbind** process when it connects to a domain controller. As result, the cache is now updated as expected.

Enhancement**BZ#915455**

The **smbd** daemon expected the old printing databases of Samba 3.5 to be in the UTF-8 format. However, the databases could be also in a different format, for example in Latin-1. Consequently, **smbd** could not migrate the database in this case. This update enhances the **net** utility, which is used for administration of Samba and remote CIFS servers, to be able to encode the database correctly and convert it to UTF-8. As a result, **smbd** can now migrate the databases as expected.

Users of samba are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add this enhancement. After installing this update, the smb service will be restarted automatically.

8.190. SAMBA4

8.190.1. RHSA-2013:1543 — Moderate: samba4 security and bug fix update

Updated samba4 packages that fix one security issue and two bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

Samba is an open-source implementation of the Server Message Block (SMB) or Common Internet File System (CIFS) protocol, which allows PC-compatible machines to share files, printers, and other information.

Security Fix

CVE-2013-4124

An integer overflow flaw was found in the way Samba handled an Extended Attribute (EA) list provided by a client. A malicious client could send a specially crafted EA list that triggered an overflow, causing the server to loop and reprocess the list using an excessive amount of memory.

Note: This issue did not affect the default configuration of the Samba server.

Bug Fixes

BZ#882338

When Samba was installed in the build root directory, the RPM target might not have existed. Consequently, the find-debuginfo.sh script did not create symbolic links for the libwbclient.so.debug module associated with the target. With this update, the paths to the symbolic links are relative so that the symbolic links are now created correctly.

BZ#911264

Previously, the samba4 packages were missing a dependency for the libreplace.so module which could lead to installation failures. With this update, the missing dependency has been added to the dependency list of the samba4 packages and installation now proceeds as expected.


All samba4 users are advised to upgrade to these updated packages, which contain a backported patch to correct these issues.

8.191. SANLOCK

8.191.1. RHBA-2013:1632 — sanlock bug fix and enhancement update

Updated sanlock packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The sanlock packages provide a shared storage lock manager. Hosts with shared access to a block device or a file can use sanlock to synchronize their activities. VDSM and libvirt use sanlock to synchronize access to virtual machine images.

**NOTE**

The sanlock packages have been upgraded to upstream version 2.8, which provides a number of bug fixes and enhancements over the previous version, including a new API provided for applications to request the release of a resource. (BZ#960989)

Bug Fix**BZ#961032**

Previously, the wdmd daemon did not always select the functional device when some watchdog modules provided two devices. Consequently, wdmd did not work correctly in some instances. A patch has been applied to address this bug, and wdmd now verifies the state of both devices and selects the one that works properly.

Enhancements**BZ#960993**

With this update, a new API has been provided for applications to verify the status of hosts in a lockspace.

BZ#966088

With this update, a new API has been provided for applications to check the status of hosts that are holding a resource lease.

Users of sanlock are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.192. SBLIM-CMPI-FSVOL**8.192.1. RHBA-2013:1691 — sblim-cmpi-fsvol bug fix update**

Updated sblim-cmpi-fsvol packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The sblim-cmpi-fsvol package provides the filesystem and volume management instrumentation allowing users to obtain information about mounted and unmounted file systems by use of CIM or WBEM technology and infrastructure.

Bug Fix**BZ#921482**

While enumerating the CIM_UnixLocalFileSystem class instances, certain file systems were shown as disabled although they were mounted. This happened when the entries in the /etc/fstab/ directory used symbolic links or Universally Unique Identifiers (UUIDs). A patch has been provided to fix this bug and the mounted file systems are now displayed correctly.

Users of sblim-cmpi-fsvol are advised to upgrade to these updated packages, which fix this bug.

8.193. SBLIM-SFCC

8.193.1. RHBA-2013:1692 — sblim-sfcc bug fix update

Updated sblim-sfcc packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The small footprint CIM client library (sblim-sfcc) is a C API allowing client applications to interface with CIM (Common Information Model) implementations (e.g. CIM servers). Due to its small memory and disk footprint it is well-suited for embedded environments.

Bug Fix

BZ#875011

Previously, when the attribute "EnumerateInstances" was present in a KEYVALUE pair, the sfcc client failed to parse the XML file and the client terminated unexpectedly with a core dump being generated. This bug has been fixed and the sfcc client now successfully parses XML files with the aforementioned attribute.

Users of sblim-sfcc are advised to upgrade to these updated packages, which fix this bug.

8.194. SBLIM-WBEMCLI

8.194.1. RHBA-2013:1047 — sblim-wbemcli bug fix update

Updated sblim-wbemcli packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

WBEM (Web-Based Enterprise Management) CLI is a standalone, command-line WBEM client. It can be used in scripts for basic system management tasks.

Bug Fixes

BZ#745264

Previously, the spec file of the sblim-wbemcli package contained a requirement for the tog-pegasus CIM server, which could cause problems. This update amends the spec file and top-pegasus is no longer required when installing sblim-wbemcli.

BZ#868905

Due to incorrect usage of the curl API in the code, when the wbemcli utility was called with an HTTPS scheme, wbemcli terminated unexpectedly with a segmentation fault. With this update, the curl API is used properly and wbemcli no longer crashes in the described scenario.

Users of sblim-wbemcli are advised to upgrade to these updated packages, which fix these bugs.

8.195. SCL-UTILS

8.195.1. RHBA-2013:1554 — scl-utils bug fix update

Updated scl-utils packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The scl-utils packages provide a runtime utility and RPM packaging macros for packaging Software Collections. Software Collections allow users to concurrently install multiple versions of the same RPM packages on the system. Using the scl utility, users may enable specific versions of RPMs, which are installed into the /opt directory.

Bug Fixes

BZ#949995

Previously, detection of collections that were specified to be enabled was done in the incorrect place in the code. Thus, if users wanted to enable multiple collections with a single command, only the first one was enabled, while the rest were ignored. With this update, the package scans all the arguments now, as opposed to the original approach where only the first one was taken, and all specified collections are now enabled.

BZ#955669

When starting an inspection of what collections had already been enabled, an incorrect variable was taken as a source of this information. In a specific case, when users ran a shell in an scl-enabled environment, and tried to enable an already-enabled collection, then the collection was enabled twice. This could have led to errors in a newly created environment, and, consequently, could have caused problems with applications running in this environment. This update accepts the correct variable as a source of information about the already-enabled collections, and the collections are no longer enabled multiple times.

BZ#957185

Previously, python27 required a specific byte compiler, thus, the build of python27 collection failed, because it used the incorrect byte compiler. With this update, python27 utilizes a new function to override various RPM macros and, thus, can be compiled.

BZ#957754

If the PATH variable was not set as scl-utils expected it, executing the "scl enable" command produced a "command not found" error. This was caused by the scl utility calling the scl_enabled command without an absolute path, and relying on the PATH that the user has set. This update uses an absolute path when calling the scl_enabled helper script, thus it does not rely on PATH any more, and the aforementioned errors no longer occur.

BZ#964058

Previously, when enabling collections, a check whether a collection was enabled was performed. However, independently of the result the collection was always enabled. Consequently, if a single collection had been stated multiple times on the command line, it was enabled multiple times. In the case of some destructive enable scriptlets, this might have led to unexpected behavior. This update runs the enable scriptlet only if the collection has not been enabled before, and attempts to enable a collection multiple times in one environment will be ignored.

Users of scl-utils are advised to upgrade to these updated packages, which fix these bugs.

8.196. SCSI-TARGET-UTILS

8.196.1. RHBA-2013:1684 — scsi-target-utils bug fix update

Updated scsi-target-utils packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The scsi-target-utils packages contain a daemon and tools to setup Small Computer System Interface (SCSI) targets. Currently, software Internet SCSI (iSCSI) and iSCSI Extensions for RDMA (iSER) targets are supported.

Bug Fixes

BZ#910638

Previously, the `tgtadm` utility did not check for the presence of the `libaio` library to enable the asynchronous I/O types of backend storage. Consequently, attempts to add a new iSCSI target device with the `"tgtadm --bstype aio"` command failed with an "invalid request" error message. This update adds `libaio` as a runtime dependency. Now, using the `"--bstype aio"` option with the `tgtadm` utility no longer fails and attempts to add a new logical unit work as expected.

BZ#813636

Prior to this update, when interruptions were occurring in the network, then reconnection of the TCP protocol did not work properly. As a consequence, memory leaks occurred in the `tgtd` daemon under these circumstances. This bug has been fixed and the TCP reconnection now works correctly in the described scenario.

BZ#865739

Previously, the `tgtd` daemon did not report its exported targets properly if configured to report them to an Internet Storage Name Service (iSNS) server. Consequently, running the `"iscsiadm -m discoverydb -t isns"` command failed. This bug has been fixed and `tgtd` now reports its exported targets correctly in the described scenario.

BZ#922270

Previously, it was not possible to supply command-line parameters to the `tgtd` daemon. With this update, it is possible to set the `TGTD_OPTIONS` variable containing the parameters and use it in the `/etc/sysconfig/tgtd` file.

Users of `scsi-target-utils` are advised to upgrade to these updated packages, which fix these bugs. All running `scsi-target-utils` services must be restarted for the update to take effect.

8.197. SEABIOS

8.197.1. RHBA-2013:1655 — seabios bug fix and enhancement update

Updated `seabios` packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The `seabios` packages contain an open-source legacy BIOS implementation which can be used as a coreboot payload. It implements the standard BIOS calling interfaces that a typical x86 proprietary BIOS implements.

Bug Fixes

BZ#846519

Due to a bug in the Advanced Configuration and Power Interface (ACPI) description table, a stop error (known as Blue Screen of Death, or BSoD) could occur on certain Windows guests. The problem was observed on the guests using the `virtio-win` small computer system interface (SCSI) drivers during heavy S3 and S4 power state transitions, for example when running the `CrystalDiskMark` benchmark. A patch has been applied to fix this bug, and a stop error no longer occurs in the described scenario.

BZ#846912

The multiple ACPI description table (MADT) previously contained an incorrect definition. Consequently, after switching to the S3 or S4 power state, the SCSI driver could not have been disabled and was always attached to the Global System Interrupt (GSI) number 9. The incorrect MADT definition has been fixed and the SCSI driver can now be disabled and enabled in this situation as expected.

BZ#888633

The SeaBIOS utility previously allowed booting even from deselected devices when no bootable device was found. This problem has been fixed by adding support for the HALT instruction, which prevents SeaBIOS from default boot attempts if no bootable device is available.

Enhancements**BZ#876250**

With this update, the SMBIOS GUID number (the same number as the system's UUID) is now displayed on the BIOS Power-on self-test (POST) screen.

BZ#963312

This update modifies the virsh management user interface so the "virsh dump" command now supports automatic core dumps. With appropriate kdump mechanism settings, the vmcore file is automatically captured and the subsequent actions are executed automatically.

Users of seabios are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.198. SELINUX-POLICY

8.198.1. RHBA-2013:1598 — selinux-policy bug fix and enhancement update

Updated selinux-policy packages that fix a number of bug fixes and add various enhancements are now available for Red Hat Enterprise Linux 6.

The selinux-policy packages contain the rules that govern how confined processes run on the system.

Bug Fixes**BZ#872542**

When SELinux was in enforcing mode and the **AWStats** utility was configured to purge **httpd** log files, AVC messages were generated due to missing SELinux policy rules for this setup. To fix this bug, the **awstats_purge_apache_log_files** Boolean was added. When enabled, the Boolean allows **AWStats** to purge the log files. Thus, the AVC messages are no longer returned.

BZ#878148

Due to a missing SELinux policy rule, the **httpd** daemon did not have permissions for searching the **/var/lib/cobbler/webui_sessions/** directory. Consequently, the user was not able to log into the Cobbler Web User Interface (UI). With this update, the SELinux policy has been updated and the user is now able to use the Cobbler Web UI as expected.

BZ#890646, BZ#890647, BZ#892024

When SELinux was in enforcing mode, the following problems related to the **postfix** service occurred:

- The **postfix** service was unable to connect to the MySQL database.
- The **sysadm_u** SELinux user was not able to execute the **postqueue -p** command correctly.
- The **postfix** daemon was not able to list the content of the **/tmp/** directory.
- When the Sender Policy Framework (SPF) verification was enabled on a gateway, the **postfix-master** binary was not able to execute the **postfix-policyd-spf-perl** Postfix server.

With this update, a set of new SELinux policy rules has been added to the SELinux policy to fix these bugs. As a result, **postfix** now works as expected in the described scenarios.

BZ#903371

Previously, a proper security context for the **/usr/local/bin/x11vnc** file was missing. Consequently, SELinux in enforcing mode blocked the GNOME Display Manager (GDM) and the X.Org implementation of the X Window System from executing the **x11vnc** server utility. The **xserver_exec_t** security context for the file has been added to the SELinux policy and GDM and X.Org now work correctly in the described scenario.

BZ#906346

Due to missing SELinux policy rules, the **sysstat** utility was unable to write a device label when generating data for the **sar** command. With this update, the SELinux policy has been updated to allow **sysstat** to work correctly.

BZ#906773

Previously, a proper security context for the **/bin/yum-builddep** file was missing. Consequently, SELinux in enforcing mode returned an error after installation of the sendmail package using the **yum-builddep** command. The security context has been updated to **rpm_exec_t** and the installation using **yum-builddep** now proceeds as expected.

BZ#908095

Due to incorrect SELinux policy rules, an attempt to use the **df_inode** plug-in of the **Munin** utility caused AVC messages to be returned. The policy rules have been updated and the plug-in now works as expected.

BZ#909857, BZ#983601, BZ#1003571, BZ#1021566

When SELinux was in enforcing mode, the following problems related to the **tgttd** daemon occurred due to insufficient SELinux policy rules:

- The **tgttd** daemon was not able to connect to the TCP port 3205 when it was running on a server together with the **iSNSd** daemon. Consequently, **tgttd** failed to discover the Internet Storage Name Service (iSNS) target.
- The **tgttd** daemon failed to access the **/dev/infiniband/uverbs0** device due to missing SELinux labeling for the device.
- The **SYS_RAWIO**, **SYS_ADMIN** and **IPC_LOCK** capabilities were missing.

- The **tgtd** daemon failed to access the **/dev/sg0** device.

The appropriate SELinux policy rules have been added to fix these bugs and **tgtd** now works as expected in the described scenarios.

BZ#912295

Previously, when multiple devices were added to the system, a **udev** rule restarted the **ktune** services for each new device. This could lead to many restarts in a short period of time. The multiple restarts could trigger a race condition in the kernel, which cannot be currently fixed. The **tuned** daemon code has been modified not to trigger more than one restart per 10 seconds, thus preventing the race condition from occurring.

BZ#913673

When the **cgrulesengd** daemon attempted to use the **inotifyfs** scripts for monitoring file-system changes, SELinux denied the daemon to access to the scripts due to the insufficient SELinux policy. This update adds a new SELinux policy rule to fix this bug and **cgrulesengd** can now use **inotifyfs** as expected.

BZ#915729, BZ#966203, BZ#984903

When SELinux was in enforcing mode, the following problems related to the **system-config-kdump** utility occurred due to insufficient SELinux policy rules:

- The **kexec** feature running in the **kdumpgui_t** SELinux domain was not able to access the **kcore** file.
- The **system-config-kdump** was unable to write to the **/boot/efi/EFI/redhat/grub.cfg** file.
- The **system-config-kdump** failed to write the **zipl** information.

The appropriate SELinux policy rules have been added to fix these bugs and **system-config-kdump** now works as expected.

BZ#917157, BZ#991024

Previously, Nagios Remote Plugin Executor (NRPE) was not allowed to execute the **sudo** utility due to missing SELinux policy rules. Consequently, when users used NRPE and their own **Nagios** plugins for monitoring servers, an attempt to call the **status** action of the **init.d** script for the supplied service, to determine the health of the service, failed. The appropriate SELinux policy rules have been updated so that NRPE can now use the **sudo** utility as expected.

BZ#919192

Due to an incorrect label of the **/var/lock/subsys/dirsrv-admin** file, an attempt to restart the Administration server using the console or the command line failed. As a consequence, AVC denial messages were returned. This update adds the proper default security context for the file and denial messages are now no longer returned.

BZ#919893

Previously, a proper security context for the **/sbin/ip6tables** file was missing. Consequently, SELinux in enforcing mode caused failures in the **Shorewall** utility. With this update, the security context has been updated to **iptables_exec_t**. As a result, **Shorewall** works as expected.

BZ#921234

Due to missing SELinux policy rules, the **abrt_t** SELinux domain was not allowed to make a transition to the **prelink_t** SELinux domain. As a consequence, the RPM verification of a package, which provided binary of a package that had terminated unexpectedly, failed during the Automatic Bug Reporting Tool (ABRT) processing. The SELinux policy has been modified to fix this bug so that the RPM verification no longer fails in the described scenario.

BZ#922028

Previously, SELinux in enforcing mode prevented the **snmptthandler** utility from performing any operations in the **/var/spool/snmptt/** directory due to the incorrect security context of the directory. With this update, the context has been updated to **snmpd_var_lib_t** so that the utility now works as expected.

BZ#922135

Due to incorrect SELinux policy rules, the **Nagios** application was unable to temporary store a file with its test results in the **/var/spool/nagios/checkresults/** directory. This update fixes the relevant SELinux policy rules and **Nagios** is no longer prevented from storing the file in this directory.

BZ#927003

The Network Information Service (NIS) master can be configured with other machines running as NIS slaves. Previously, when a NIS client changed the NIS password, a new AVC message was logged into the **/var/log/audit/audit.log** file. This was because SELinux did not allow the **yppus** utility to connect to the Transmission Control Protocol (TCP) 111 port. With this update, the appropriate SELinux policy rules have been modified and the AVC message is no longer logged in the described scenario.

BZ#927973

Due to the incorrect SELinux policy, running the Apache HTTP Server alongside with the **postfix** agent did not work correctly. As a consequence, the **postdrop** utility, which was labeled with the **httpd_t** SELinux label, was unable to access the **/var/spool/postfix/maildrop/** directory. With this update, the **httpd_can_sendmail** Boolean has been updated to allow **postdrop** to access the directory.

BZ#947772

When SELinux was in enforcing mode, the **sanlock-helper** utility was not allowed to send a SIGKILL signal to any process, which was registered to the **sanlock** daemon. The relevant SELinux policy rules have been modified with this update and **sanlock-helper** is now able to send the SIGKILL signal to the registered processes.

BZ#950103

Due to insufficient SELinux policy rules, a transition between the **pegasus_t** and the **mount_t** SELinux domains did not work correctly. Consequently, when the **OpenPegasus** Web-Based Enterprise Management (WBEM) services tried to retrieve information about a file system using the **wbemcli** utility, the access to the mount was denied by SELinux. With this update, the SELinux policy has been modified and **OpenPegasus** is now able to access the mount in the described scenario.

BZ#952621

When SELinux was in enforcing mode, the **sandbox** SELinux domains were not able to use inherited user terminals due to missing SELinux policy rules. With this update, the respective rules have been updated to allow **sandbox** domains to use these terminals.

BZ#953180

Due to insufficient SELinux policy rules, when the **s2s** service was used in the mixed Red Hat Network Satellite and Red Hat Network Satellite Proxy environment, the following AVC message was returned in the **audit.log** file:

```
type=AVC msg=audit(1364300742.715:101611): avc: denied { name_connect }
for pid=2278 comm="s2s" dest=5269
scontext=system_u:system_r:jabberd_t:s0
tcontext=system_u:object_r:jabber_interserver_port_t:s0
tclass=tcp_socket
```

The appropriate SELinux rules have been added to fix this bug and the AVC message is no longer returned in such a case.

BZ#956720

Previously the **opasswd** and the **opasswd.old** files were labeled with the **etc_t** SELinux context. However, these files included sensitive information and were supposed to be labeled with the **shadow_t** context. With this update, the SELinux policy has been modified and the files are now correctly labeled with **shadow_t** as expected.

BZ#957012

Previously, clock devices (**/dev/ptp***) were incorrectly labeled with the **device_t** SELinux label instead of **clock_device_t**. This update provides a patch to fix this bug and the clock devices are now correctly labeled.

BZ#957023

Previously, SELinux in enforcing mode prevented the **svnserv** daemon from using the TCP port 3690. The appropriate SELinux policy rules have been updated and **svnserv** can now use the port as expected.

BZ#957265

Due to missing SELinux rules, a transition between the **aide_t** and the **prelink_t** SELinux domains was not possible. As a consequence, when SELinux was running in enforcing mode, the **aide --check** command executed inside a **cron** job did not work correctly. The respective SELinux rules have been updated to fix this bug and the command now works as expected.

BZ#958682, BZ#975921, BZ#1009449

Previously, the **mysqld_safe** script was unable to execute a shell (**/bin/sh**) with the **shell_exec_t** SELinux security context. Consequently, the **mysql55** and **mariadb55** Software Collection packages were not working correctly. With this update, SELinux policy rules have been updated and these packages now work as expected. In addition, the **mysqld_safe** SELinux policy has been modified to allow the **SYS_NICE** capability.

BZ#966106

When using certain versions of the **Quantum** service with **netns** support, SELinux denied various operations, which caused **Quantum** to terminate unexpectedly. Moreover, due to a “dontaudit” rule for

the operations, AVC messages were not returned unless SELinux was running in permissive mode. The appropriate SELinux policy has been fixed so that SELinux no longer denies the operations and **Quantum** failures no longer occur in the described scenario.

BZ#966515

Previously, enabling the **ftp_homdedir** Boolean allowed certain rules, that were not supposed to be allowed by the Boolean. The relevant SELinux policy has been modified and the Boolean now allows only the rules that it is supposed to.

BZ#966635

Previously, the **Munin** Common Gateway Interface (CGI) scripts was labeled incorrectly, and therefore ran in an incorrect SELinux domain. The file context for the scripts has been updated to **httpd_munin_script_exec_t** and the scripts now run in the correct SELinux domain.

BZ#966640

Previously, the **/var/log/syslog-ng** file was incorrectly labeled with the **syslog_var_run_t** SELinux security context. Consequently, when SELinux was running in enforcing mode, the **logwatch** utility was unable to access the file. With this update, the security context for the **syslog-ng** file has been modified to **var_log_t** and **logwatch** can now access the file as expected.

BZ#971594

Previously, an attempt to attach a Logical Volume Management (LVM) volume to a Red Hat OpenStack 3 instance failed due to the incorrect SELinux policy and AVC denial messages were returned. The relevant SELinux policy rules have been modified to add an additional Multi-Category Security (MCS) attribute for the **hald_t** SELinux domain. As a result, the AVC denial messages are now no longer returned in the described scenario.

BZ#973156

Previously, the **/etc/yaboot.conf** file was incorrectly labeled with the **etc_t** SELinux security context. With this update, the security context has been changed to the **bootloader_etc_t**.

BZ#974932

The **SETUID** and **SETGID** capabilities were missing in the SELinux policy. As a consequence, when SELinux was in enforcing mode, the **rsyslog** utility was unable to drop privileges with the **\$PrivDropToUser** and **\$PrivDropToGroup** options. With this update, the missing capabilities have been added to the SELinux policy and **rsyslog** can now drop privileges as expected.

BZ#978993

Due to incorrect SELinux policy rules, SELinux prevented the **chronyd** daemon from using the **SYS_NICE** capability. The capability is required by the **sched_setscheduler()** function. With this update, the SELinux policy rules has been modified to allow the daemon to use **SYS_NICE**.

BZ#983217

Previously, a transition from the **dovecot_t** SELinux domain to the **oddjob_mkhome_dir_t** SELinux domain was not allowed. Consequently, an attempt to create a user home directory alongside with the Dovecot server and the **pam_oddjob_mkhome_dir** module enabled failed and AVC messages were returned. The SELinux policy has been modified so that the transition is now allowed.

BZ#995434

SELinux running in enforcing mode prevented the **lldpad** service from communicating with the **fcoemon** service. As a consequence, the user was not able to create a virtual machine in Virtual Machine Manager (**virt-manager**) and the following AVC message was returned:

```
type=AVC msg=audit(1376046443.294:69876): avc: denied { sendto } for
pid=2755 comm="lldpad" path=003030303232
scontext=system_u:system_r:lldpad_t:s0
tcontext=system_u:system_r:fcoemon_t:s0 tclass=unix_dgram_socket
```

The appropriate SELinux policy has been fixed and users are now able to create virtual machines as expected.

BZ#998663

Previously, the SELinux policy prevented running virtual machines based on volumes located in the **/var/run/vdsm/storage/** VDSM's daemon directory. As a consequence, an attempt to run such a virtual machine terminated unexpectedly with an error. With this update, the **svirt_t** SELinux domain has been updated to read symbolic links in the **/var/run/** directory. As a result, the virtual machines no longer fail in the described scenario.

BZ#1005196, BZ#1005250

Due to incorrect SELinux policy rules, certain SELinux domains were unable to access the **/sys/devices/system/cpu/** directory. Consequently, such domains could not get information from the directory. With this update, the relevant SELinux policy rules have been updated to allow the domains access to the **/sys/devices/system/cpu/** directory.

BZ#1005806

With the Multi-Level Security (MLS) SELinux policy enabled, the **xinetd** daemon failed to execute a shell script and the following error message was returned:

```
xinetd[2771]: execv( /usr/local/ea14_testing/audit-test/utils/network-
server/pidfile_kill.sh ) failed: Permission denied (errno = 13)
```

The appropriate SELinux rules have been updated to allow **xinetd** to execute shell scripts.

BZ#1006952

Due to insufficient SELinux policy rules, an attempt to start a QEMU process using the **libvirt** library failed with an error. With this update, the SELinux policy has been modified and QEMU processes now start as expected.

BZ#1009661

Due to insufficient SELinux policy rules, the **beaker** jobs failed during automatic wireless testing and an AVC denied message was returned. Consequently, users were unable to use the wireless connection. The appropriate SELinux policy rules have been updated to fix this bug so that users can now use the wireless connection in the described scenario.

BZ#1009838

Due to missing SELinux policy rules, when the system was set up to use the **yppasswdd** daemon on a server, the **rpc.yppasswdd** binary was now allowed to read the **/var/run/utmp** file and list the content of the **/boot/** directory. The relevant SELinux policy has been updated and the daemon can

now access the **utmp** file and the **/boot/** directory as expected.

BZ#1009859

When the system was set up to the Concurrent Versions System (CVS) server using Pluggable Authentication Module (PAM) for client authentication, the CVS binary was not allowed to read the **/var/run/utmp** file. This update fixes the relevant SELinux policy to allow CVS to read the file as expected.

Enhancements**BZ#926022**

With this enhancement, a new Boolean, **ftpd_use_fusefs**, has been added to the SELinux policy. When enabled, this Boolean allows the GlusterFS mounts to be used for the File Transfer Protocol (FTP) data directory.

BZ#854963, BZ#876334, BZ#881834, Bz#891779, BZ#1000521

The **pand**, **haproxy**, **watchdog**, **lldpad**, and **openhpid** daemons ran in the **initrc_t** SELinux domain. With this enhancement, SELinux support has been added for the daemons and they now use their own separate SELinux domains.

BZ#871437

With this enhancement, a new SELinux policy for the **smstools** package is provided.

BZ#880728, BZ#986198

Previously, the manual pages did not include all updated SELinux policy rules. With this update, the actual SELinux policy is included in the **selinux-policy** package. As a result, such manual pages are up-to-date.

BZ#889120, BZ#915151, BZ#923246, BZ#924843, BZ#1011963,

Previously, the **pacemaker** resource manager did not have its own SELinux policy defined and used the **initrc_t** domain. With this update, all cluster administrative services including **pacemaker** have been merged together to the **cluster_t** SELinux domain. In addition to this merge, all other Red Hat Cluster services have been updated to use the **cluster_t** domain.

BZ#859651, BZ#1004380, BZ#1010324

The **git_shell_t** SELinux type has been removed from the SELinux policy. With this enhancement, the updated SELinux policy for the **Git** control system is provided.

BZ#890554

With this enhancement, the SELinux policy for the **Zabbix** monitoring system has been updated.

BZ#915314

With this enhancement, a set of new rules, which allows the user to mount the Gluster file system, has been added to the SELinux policy.

BZ#922732, BZ#966387

A new SELinux file type and label has been added for the **/var/lib/openvpn/** directory. In addition, the SELinux policy has been updated to allow OpenVPN to manage its own log files.

BZ#928020, BZ#955189, BZ#979421, BZ#999471, BZ#1002593

With this enhancement, the **amavis_t**, **clamd_t**, **clamscan_t**, **freshclam_t** SELinux domains have been merged to the **antivirus_t** SELinux domain.

BZ#952827

With this update, SELinux support for 27017, 28017, 27018, 28018, 27019 and 28019 ports has been added. These now ports use their separate **mongod_port_t** SELinux port type.

BZ#953652, BZ#963465, BZ#968344, BZ#969485

With this update, the SELinux policy for the **OpenShift** application platform has been updated to reflect the latest upstream policy.

BZ#953754

The file contexts for all **Nagios** plug-ins located in the **usr/lib(64)?/nagios/plugins/** directory have been updated to the **nagios_unconfined_plugin_exec_t** context.

BZ#955774

With this enhancement, two new Booleans have been added to the SELinux policy. The **tftp_use_nfs** Boolean allows The Trivial File Transfer Protocol (TFTP) to read from NFS volumes for public file transfer services. The **tftp_use_cifs** Boolean allows TFTP to read from CIFS volumes.

BZ#959554

The new Shared System Certificates feature has added new locations, from which system trusted certificated and blacklist information could be read. With this enhancement, SELinux file contexts have been updated accordingly.

BZ#964345

The SELinux policy related to the QEMU Guest Agent (**qemu-ga**) has been updated according to new **qemu-ga** features and functionality.

BZ#968403

With this update, the SELinux policy for the Oracle Automatic Storage Management (ASM) has been updated to reflect the latest upstream policy.

BZ#977047

The Zettabyte File System (ZFS) has been added to the **xattr** list of supported file systems. With this enhancement, the SELinux policy has been updated accordingly.

BZ#979432

The new **openvpn_run_unconfined** Boolean has been added to the SELinux policy. When enabled, the Boolean allows OpenVPN to execute unconfined scripts.

BZ#986883

With this update, the SELinux policy for Internet Protocol Security (IPsec) has been updated to reflect the latest upstream policy.

BZ#1006370

With this update, the prefix of the **openstack-selinux** policies has been changed from “quantum” to “neutron”.

BZ#1011973

With this enhancement, the TCP port 9000 is labeled with the **httpd_port_t** SELinux label.

Users of selinux-policy are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.199. SETUPTOOL

8.199.1. RHBA-2013:0891 — [setuptools bug fix update](#)

Updated setuptools packages that fix one bug are now available for Red Hat Enterprise Linux 6.

Setuptools is a user-friendly text mode menu utility which allows the user to access all of the text mode configuration programs included in the operating system distribution.

Bug Fix

BZ#883581

The `/usr/share/man/man1/setup.1.gz` file no longer has the executable flags set.

Users of setuptools are advised to upgrade to these updated packages, which fix this bug.

8.200. SG3_UTILS

8.200.1. RHBA-2013:0956 — [sg3_utils bug fix update](#)

Updated sg3_utils packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The sg3_utils packages contain a collection of tools for SCSI devices that use the Linux SCSI generic (sg) interface. This collection includes utilities for database copying based on “dd” syntax and semantics (the “sg_dd”, “sgp_dd” and “sgm_dd” commands), INQUIRY data checking and associated pages (“sg_inq”), mode and log page checking (“sg_modes” and “sg_logs”), disk spinning (“sg_start”) and self-tests (“sg_senddiag”), as well as other utilities. It also contains the `rescan-scsi-bus.sh` script.

Bug Fix

BZ#920687

The logic used in the `show_devices()` function (as used by the “sginfo” command) to identify the available storage devices opened these devices by scanning the `/dev` directory. Consequently, the `/dev/snapshot` file was opened, causing the system to prepare for suspend/hibernate and, as a part of that, to block hot-added CPUs from being activated. To fix this bug, the logic used in the `show_devices()` function to decide, which devices to open, has been adjusted. As a result, running the “sginfo -l” command no longer has the unintended side effect of blocking the activation of hot-added CPUs.

Users of sg3_utils are advised to upgrade to these updated packages, which fix this bug.

8.201. SLAPI-NIS

8.201.1. RHBA-2013:1671 — slapi-nis bug fix update

Updated slapi-nis packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The slapi-nis package contains the NIS server plug-in and the Schema Compatibility plug-in for use with the 389 directory server.

Bug Fixes

BZ#923336

Due to a bug in the NIS Server plug-in, when multiple clients were connected to the server over stream sockets, some of these clients were disconnected earlier than expected. This bug has been fixed and clients are no longer disconnected too early in the described case.

BZ#967468

Prior to this update, the NIS Server plug-in did not release memory that was used to hold a decoded client request. Consequently, the directory server used an excessive amount of memory as it processed more requests from NIS clients. The bug has been fixed and the directory server now consumes an expected amount of memory.

Users of slapi-nis are advised to upgrade to these updated packages, which fix these bugs.

8.202. SOS

8.202.1. RHBA-2013:1688 — sos bug fix and enhancement update

An updated sos package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The sos package contains a set of tools that gather information from system hardware, logs and configuration files. The information can then be used for diagnostic purposes and debugging.

Bug Fixes

BZ#876309

The SELinux plug-in used some commands that are obsolete on modern Linux distributions so the sosreport utility was unable to collect some information from the SELinux tools and diagnostics. The plug-in has been updated to reflect changes in the SELinux tools and diagnostics so that sosreport is now able to collect more information from these components.

BZ#883811

Previous versions of sos did not mask passwords in libvirt's XML configuration files and output of the corosync-obctl command so the passwords may have been disclosed to recipients of sosreport data. This update modifies the respective libvirt and corosync plug-ins so that passwords are now left out when collecting sos data from the aforementioned sources.

BZ#888488

Previously, when executing external commands, sos always used the user's environment settings unless the environment was explicitly specified by a plug-in used for collecting sos data.

Consequently, the collected output was a subject to locale and custom settings of the user running the `sosreport` command, which could be undesirable for further processing of the data. With this update, `sos` runs all external commands with a consistent `LC_ALL` setting and the command output is now collected using the `C` locale.

BZ#888589

The `sosreport` utility previously verified all installed packages by default, which was highly demanding on CPU and memory usage. To avoid this situation, the `rpm` plug-in now contains a fixed list of packages to verify, including core system packages such as the kernel packages.

BZ#888724

Previous versions of `sos` did not preserve the permissions of the collected files. File permissions in `sosreport` archives could have been inconsistent with file permissions on the host system, potentially misleading the user. With this update, `sos` preserves ownership and permissions of the collected files.

BZ#913201

The `sosreport` utility previously could cause unexpected RPC failures on the local system by attempting to copy RPC channel files from the `proc` file system (`/proc/net/rpc/*/channel`). These files are now blacklisted from collection and the `sosreport` command can no longer interfere with active RPC communications.

BZ#924925

The `openswan` plug-in previously collected output of the `"ipsec barf"` command to obtain VPN related diagnostic information. This could cause `sosreport` to appear unresponsive when running on systems that contained accounts with large UIDs and had installed a version of `openswan` affected by bug 771612. With this update, the `ipsec barf` command is no longer run by default, and the problem can no longer occur in this scenario, unless the `barf` functionality is explicitly enabled from the command line.

BZ#947424

The `devicemapper` plug-in used an obsolete syntax to obtain information from the `udev` subsystem. The plug-in called the `"udevinfo"` command instead of the actual command, `"udevadm info"`. This has been fixed with this update, and the correct property data can now be collected for the relevant block device types.

BZ#966602

The `sosreport` command incorrectly assumed that the `tar` program would always write data on standard output by default. Consequently, when the `TAPE` environment variable was set, data may have been unexpectedly written to a tape device, or another location expanded to by this variable. The `sosreport` has been modified to always call the `tar` command with the `"-f"` option, forcing data to be written to standard output. Users who set the `TAPE` variable in their environment can run `sosreport` without risking that data on existing tape devices could be overwritten.

BZ#986301

Previous versions of `sos` allowed passwords from `luci` configuration files to be collected by the `cluster` module so the passwords may have been disclosed to recipients of `sosreport` data. This update modifies the `cluster` module so that `luci` passwords are now left out from the collected data.

BZ#986973

Previous versions of the `sos` package called the `"wbinfo -u"` command to collect user information from domains visible to the system `Winbind` configuration. However, the `wbinfo` command may have used

very large amounts of memory and CPU time on large Active Directory installations with many trusted domains. As a consequence, sosreport appeared to be unresponsive and may have triggered out-of-memory conditions for other processes. The sosreport command has been modified to use the "--domain=.'" switch with the wbinfo command, which restricts data collection to the local domain. The problem no longer occurs in the described scenario.

BZ#987103

Previous versions of sos collected the file /etc/krb5.keytab on systems where kerberos authentication is configured. This file contains encrypted keys and is of limited diagnostic value. A summary of entries in the file is now obtained using the klist command instead.

Enhancements

BZ#868711

The output of the "gluster volume geo-replication-status" command may be important for debugging problems related to Gluster geographic replication. Therefore, the gluster plug-in now collects this diagnostic output by default.

BZ#907861

The ID mapping daemon (idmapd) controls identity mappings used by NFS services and may be important for diagnostic and troubleshooting efforts. Therefore, the idmad.conf configuration file is now collected on NFS client and server hosts, and can be analyzed in the sosreport utility.

BZ#924338

The sosreport utility now allows collecting configuration files for the Open Hardware Platform Interface (OpenHPI) components.

BZ#924839

The sosreport utility now collects kernel log data (dmesg logs) from vmcore dump files that are found on the system.

BZ#989292

The sos package now supports collecting of unified cluster diagnostic data with the crm_report tool.


Users of sos are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

8.203. SPICE-GTK

8.203.1. RHBA-2013:1577 — spice-gtk bug fix and enhancement update

Updated spice-gtk packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The spice-gtk packages provide a **GIMP Toolkit (GTK+)** widget for **SPICE** (Simple Protocol for Independent Computing Environments) clients. Both Virtual Machine Manager and Virtual Machine Viewer can make use of this widget to access virtual machines using the **SPICE** protocol.



NOTE

The `spice-gtk` packages have been upgraded to upstream version 0.19, which provides a number of bug fixes and enhancements over the previous version, including support tunneling through the HTTP proxy server, improved multi-monitor support, and various USB redirection fixes. (BZ#[961452](#))

Bug Fixes

BZ#[980400](#)

The `polkit` utility is built against newer `GTK+` and `GLib` versions, thus it has a runtime dependency on these versions. Previously, upgrading `spice-gtk` without upgrading `GTK+` and `GLib` at the same time caused applications using `polkit` to terminate unexpectedly on startup. With this update, the RPM dependencies have been adjusted so that `spice-gtk` RPMs require new enough versions of `GTK+` and `GLib`. As a result, `spice-gtk` cannot be installed unless the `GTK+` and `GLib` versions it requires are installed as well.

BZ#[879352](#)

Prior to this update, `spice-gtk` connected to the server plain port by default and succeeded only if the server provided the port. However, this prevented `spice-gtk` from connecting to a secure port by default. With this update, `spice-gtk` can connect to secure port instead of always trying plain ports first.

BZ#[906558](#)

Previously, if the combination of `Shift+CTRL+V` keys was pressed on a message window in Microsoft Outlook, the `spice-gtk` client terminated unexpectedly. To fix this bug, cache palettes of unrendered bitmaps have been applied and the client no longer crashes in the aforementioned scenario.

BZ#[998529](#)

When the mouse pointer was placed over a PuTTY session with Microsoft Windows guests, the mouse pointer appeared black. The cursor contrast has been improved and the cursor is now clearly visible when hovering over PuTTY sessions.

BZ#[885101](#)

When `spice-gtk` was connecting to an unreachable host, a connection timeout error took about 2 minutes to occur. With this update, `spice-gtk` waits for 10 seconds only before reporting an unreachable host error.

BZ#[815639](#)

Previously, `spice-gtk` did not handle correctly an indication that software `Smartcard` support had already been initialized. Consequently, software `Smartcard` support stopped working after migration or restarting a guest. As a workaround, do not disable software `Smartcard` support at `spice-gtk` connection time if `libccard` reports that software `Smartcard` support is already initialized. Pursuing this workaround, software `Smartcard` support keeps working across guest reboots or migrations.

Enhancements

BZ#[948618](#)

HTTP proxy server support is now available for **SPICE** clients. The **SPICE** client now establishes the connection to the remote server by the proxy server specified by the environment `SPICE_PROXY=host:port` variable, or by the controller.

BZ#752350

Red Hat Enterprise Linux and Microsoft Windows operating systems use different end of line sequences. With this update, a new feature has been implemented, which translates end of line sequences to the target operating system during copy and paste, if the **SPICE** guest agent has support for this feature.

BZ#978405

Due to high latency and limited bandwidth, a video could become very poorly watchable. With this update, adaptive video streaming with adjusted bandwidth and latency has been implemented, which allows improved video experience over low speed networks.

Users of `spice-gtk` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.204. SPICE-PROTOCOL

8.204.1. RHEA-2013:1573 — spice-protocol enhancement update

Updated `spice-protocol` packages that add various enhancements are now available for Red Hat Enterprise Linux 6.

The `spice-protocol` packages provide header files to describe the SPICE protocol and the QXL paravirtualized graphics card. The SPICE protocol is needed to build newer versions of the `spice-client` and the `spice-server` packages.



NOTE

The `spice-protocol` packages have been upgraded to upstream version 0.12.6, which provides a number of enhancements over the previous version, including file copy support, and end-of-line conversion for copying and pasting text using clipboard when the client and guest conventions differ, or when there are different operating systems running. (BZ#965812)

Enhancement

BZ#978410

The continuity of video playbacks over limited bandwidth has improved. There are now fewer occurrences of scenes dropping and pauses, which leads to a better user experience. After this update, the video bit-rate and the playback latency are dynamically adjusted, using periodic reports from the client. The bandwidth and latency estimations of past video playbacks are now used for improving the initial parameter settings of future video playbacks.

Users of `spice-protocol` are advised to upgrade to these updated packages, which add these enhancements.

8.205. SPICE-SERVER

8.205.1. RHBA-2013:1571 — spice-server bug fix and enhancement update

Updated spice-server packages that fix a number of bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Simple Protocol for Independent Computing Environments (SPICE) is a remote display protocol for virtual environments. SPICE users can access a virtualized desktop or server from the local system or any system with network access to the server. SPICE is used in Red Hat Enterprise Linux for viewing virtualized guests running on the Kernel-based Virtual Machine (KVM) hypervisor or on Red Hat Enterprise Virtualization Hypervisors.



NOTE

The spice-server packages have been upgraded to upstream version 0.12.4, which provides a number of bug fixes and enhancements over the previous version. (BZ#[952671](#))

Bug Fixes

BZ#[823472](#)

Data accessed from the main thread, which use most SPICE channels, could be accessed by threads of other channels, such as display and cursor channels. To protect the data, an assertion check has been added to the SPICE code. However, certain calls to the sound channel interface use the Virtual CPU (vCPU) thread. Previously, these calls were rejected by the assertion check causing the SPICE server and the Kernel-based Virtual Machine (KVM) hypervisor to abort. Such calls are harmless because KVM uses global mutual exclusion (mutex) for the vCPU and I/O threads. With this update, a warning is returned instead of aborting SPICE and KVM.

BZ#[859027](#)

When the `client_migrate_info()` function was called with the `cert-host-subject` option specified and then was called without the option, on the third call, the option was freed for the second time. This was because the pointer was not set to NULL after it was first freed during the second call. This behavior caused the SPICE server to terminate unexpectedly with a segmentation fault. The underlying source code has been modified and the pointer is set to NULL when the `cert-host-subject` option is not specified. As a result, the pointer is freed only once and SPICE no longer crashes in the described scenario.

BZ#[918169](#)

When two items were to be sent to a client and the client became disconnected, the first item was cleared successfully but the second one was not. Consequently, the SPICE server terminated unexpectedly due an assertion check failure. This update applies a patch to fix this bug so that the second item is now properly cleared, too. As a result, the SPICE server no longer crashes in the described scenario.

BZ#[918472](#)

Due to a bug in the SPICE source code, an attempt to run the `getaddrinfo()` function failed with a segmentation fault. Consequently, Quick Emulator (QEMU) terminated unexpectedly. The underlying source code has been modified and QEMU no longer crashes when executing `getaddrinfo()`.

BZ#[950029](#)

When the SPICE source server was streaming video data during a migration process, the SPICE server could send stream-related messages to the SPICE client after sending a `MSG_MIGRATE`

message. This is not allowed and the client thus forwarded a wrong message instead of a **MSG_MIGRATE_DATA** message to the destination host. The destination host then aborted the migration. This update modifies the SPICE server code to ensure that only the **MSG_MIGRATE_DATA** message can be sent after sending **MSG_MIGRATE** and the migration process now successfully finish.

BZ#952666

Previously, the SPICE server did not allow creation of a surface with the “stride ≥ 0 ” path because the path was untested and it was not requested before by any QXL driver. Consequently, when a QXL driver attempted to create such a surface, SPICE terminated unexpectedly with an error on certain systems. The underlying source code has been modified to allow creation of the surface with the “stride ≥ 0 ” path. As a result, the SPICE server no longer crashes in the described scenario.

BZ#956345

Under certain circumstances, the SPICE server could abort upon a virtual machine (VM) migration. This could happen if the VM was being migrated to a new host after the previous migration to the current host within the same SPICE client session. Then, if the connection between the original host and the client was a low bandwidth connection, the new host passed an incorrect connection bandwidth value to the SPICE client causing the SPICE server to abort. This update provides a patch addressing this problem and the SPICE server now sends the correct connection bandwidth value in this scenario.

BZ#958276

Previously, the destination host did not send its multi-media time to a client during migration so that the client held the multi-media time of the source server. As a consequence, if the source and destination hosts had different multi-media time and no audio playback, video frames that were created after the migration were dropped by the client. This update applies a patch to fix this bug and video frames are no longer dropped in the described scenario.

BZ#977998

Previously, an incorrect flag that was set when sending bitmaps caused an endless loop in the client display channel. This behavior occurred especially under limited bandwidth conditions. Consequently, the SPICE server could become unresponsive. The underlying source code has been modified to fix this bug and SPICE no longer hangs in such a situation.

BZ#977998

Previously, the waiting timeout period for a client response was set to 150 seconds. This duration was too long and caused, under certain circumstances, server errors to be returned. With this update, the waiting timeout period was set to 30 seconds to prevent the server errors from occurring.

Enhancements**BZ#961848**

With this enhancement, support for the new QEMU **disable-agent-file-transfer** option has been provided. As a result, users can now filter out the file transfer messages.

BZ#978403

This update introduces the adaptive video streaming that provides better video quality. With this feature, the network bandwidth and latency are estimated and the video bit-rate and the playback latency are dynamically adjusted. Moreover, the bandwidth and latency estimations of past video playback are used for improving the initial parameters setting of future video playbacks.


All spice-server users are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.206. SPICE-VDAGENT

8.206.1. RHBA-2013:1560 — spice-vdagent bug fix and enhancement update

Updated spice-vdagent packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The spice-vdagent packages provide a SPICE agent for Linux guests.



NOTE

The spice-vdagent packages have been upgraded to upstream version 0.14.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#[951596](#))

Bug Fixes

BZ#[881020](#)

While adjusting the guest's monitor configuration, the SPICE guest agent could fail to set the resolution when it switched to full-screen mode, which left the guest's monitor configuration in an inconsistent state. This happened because the agent did not handle situations when the guest's video memory was exhausted. This update fixes the issue by reverting the guest's monitor configuration to the previous state when the agent fails to adjust the guest monitor configuration.

BZ#[894036](#)

The SPICE guest agent previously expected the guest's monitor configuration to be always continuous, and attempted to establish and maintain connections with display windows in ascending order (that is display 1, 2, 3, ...). The user was thus unable to open displays out of order and choose which display would be closed when closing a display window. The problem has been fixed by sending a sparse monitor configuration to the agent, which allows specifying of the display that is to be open or closed. Monitors to be disabled are configured as 0x0 sized monitors in the sparse monitor configuration.

BZ#[894365](#)

The user session spice-vdagent process could terminate unexpectedly when the spice-vdagentd daemon was restarted after performing certain tasks such as resizing the window, copying and pasting data and changing to full-screen mode. This happened because the SPICE guest agent attempted to free already-freed memory upon a spice-vguestd restart. The agent now frees the memory correctly, and the user session spice-vdagent process thus no longer crashes in this scenario.

BZ#[895004](#)

The SPICE guest agent for Linux logged warning messages when using multiple monitors per single qxl device, which is unnecessary since spice-vdagent supports such configurations. The spice-vdagentd daemon has been modified so the warning messages no longer occur in the log.

BZ#[999804](#)

When running the SPICE guest agent on the system without the virtio channel, the agent emitted

inadvertent error messages about a missing virtio device. This update fixes this problem by removing the respective `syslog()` call from the code and the error messages no longer occur in the system log in this situation.

BZ#1003977

When the user reconnected to a remote-viewer window after closing it while copying a large amount of data from the client to the guest, the copy-paste function stopped working. Furthermore, the user were unable to call the context menu by clicking the left mouse button. This happened because the SPICE guest agent did not release the clipboard in this situation. The agent has been modified to properly release clipboard and the problem no longer occurs.

Enhancements**BZ#799482**

The SPICE guest agent now provides support ensuring correct translation of end-of-line sequences when the client run on different operating system than is the guest's operating systems (for example, when running a Windows client and a Linux guest).

BZ#904082

This update adds support for setups with multiple X11 screens. Such a setup can be achieved using multiple `qxl` devices where each device is mapped to a separate screen. This setup brings the following limitations: only one monitor can be used per X11 screen (a `qxl` device), all monitors must have the same resolution, and resolution synchronization has to be done on guest machines since no SPICE client is present.

BZ#904084

A new `"-X"` command-line option has been added to `spice-vdagent`, which allows disabling the `ConsoleKit` framework and the `systemd-logind` service integration at runtime for setups where these services are not used.

Users of `spice-vdagent` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.207. SPICE-XPI

8.207.1. RHEA-2013:1667 — spice-xpi bug fix and enhancement update

Updated `spice-xpi` packages that fix one bug and add one enhancement are now available for Red Hat Enterprise Linux 6.

The `spice-xpi` packages provide the Simple Protocol for Independent Computing Environments (SPICE) extension for Mozilla that allows the SPICE client to be used from a web browser.

Bug Fix**BZ#882339**

Prior to this update, the `spice-xpi` browser plug-in did not remove the `/tmp/spicec-XXXXXX/spice-foreign socket` and the `/tmp/spicec-XXXXXX/` directory, so they were still present after client had exited. This bug has been fixed, and the browser plug-in now removes the above mentioned file and directory after client exits.

Enhancement

BZ#994613

Proxy support for SPICE connection has been added to the spice-xpi browser plug-in. With this update, spice-xpi is now able to pass the proxy setting to the SPICE client it spawns, for example, when opening a console from the Red Hat Enterprise Virtualization Manager portal.

Users of spice-xpi are advised to upgrade to these updated packages, which fix this bug and add this enhancement. After installing the update, Firefox must be restarted for the changes to take effect.

8.208. SSSD

8.208.1. RHBA-2013:1680 — sssd bug fix and enhancement update

Updated sssd packages that fix a number of bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The System Security Services Daemon (SSSD) provides a set of daemons to manage access to remote directories and authentication mechanisms. It provides the Name Service Switch (NSS) and the Pluggable Authentication Modules (PAM) interfaces toward the system and a pluggable back-end system to connect to multiple different account sources.

Bug Fixes

BZ#872827

In case a member of a group was outside all configured search bases, the get-group-members request could be marked as done before the caller had a chance to register a callback. As a consequence, resolving a group with members outside the search base could have appeared as stuck. The get-group-members request was fixed to call a special `tevent_req_post()` function that waits with returning the result until the caller has registered a callback. The request now works correctly even if a member is outside configured search bases.

BZ#906398

There was a `get_attribute` call used in the group processing code base that, when a nonexistent attribute was requested, could allocate an empty attribute instead of reallocating the previous attribute array. The reallocation could have invalidated existing pointers that were previously pointing to the array. In case a group contained no members at all, the array could be reallocated and existing pointers invalidated, causing the SSSD daemon to terminate unexpectedly. To fix this bug, another `get_attribute` is now used that returns the `ENOENT` error instead of creating an empty attribute. As a result, SSSD no longer crashes in the described scenario.

BZ#911329

The `pam_pwd_expiration` warning was erroneously set to the “0” value for the Kerberos provider and therefore, the password expiration warning was always displayed when the server had sent one. As a consequence, in certain environments, such as Active Directory (AD) or IPA, the warning was displayed on each login. This update applies a patch to modify this behavior and the warning is now set by default to be displayed only once in seven days.

BZ#914433

The code that created the login file for the IPA provider did not handle error conditions correctly and was unable to recover after failure of writing a SELinux label. When no `selinux-policy-targeted` directory was present on the system, the target directory that the SSSD daemon wrote to was

missing. Consequently, writing the login file failed. With this update, the underlying source code has been modified so that SSSD now correctly handles the writing failures as expected.

BZ#916997

The possibility to retrieve very large Active Directory (AD) groups instead of skipping them has been added to the previous version of Red Hat Enterprise Linux. However, this behavior could cause performance problems, because the additional resolution took a long time. To fix this bug, a new option, “`ldap_disable_range_retrieval`”, has been added allowing the SSSD daemon to skip very large AD groups.

BZ#918394

When the memory cache was reset with the `sss_cache` utility, the SSSD daemon did not close the file descriptor, which caused a file descriptor leak. The underlying source code has been modified so that the file descriptor is now closed correctly in the described scenario.

BZ#948830

Netgroups can contain nested netgroups from other sources so that the SSSD daemon resolves only one nesting level at a time and allows the glibc library to query other sources as well. However, previously, there was a full query per nesting level and therefore the nested netgroup processing was very slow. With this update, a new option, “`refresh_expired_interval`”, has been introduced. The option controls the task that updates the expired records on the background instead of waiting for the user login. As a result, the nested netgroup processing is now faster.

BZ#950874

Previously, simple access control denied access to case-insensitive domains for users that had their user names written in the uppercase characters. This update applies a patch to fix this bug so that all users are now able to log in as expected.

BZ#951086

In case the processing of an LDAP request took longer than the client timeout (60 seconds by default), upon completing the request, the PAM client could have accessed memory that was previously freed due to the client timeout being reached. As a consequence, the `sss_d_pam` process terminated unexpectedly with a segmentation fault. With this update, the SSSD daemon ignores an LDAP request result when it detects that the set timeout of this request has been reached. As a result, the `sss_d_pam` process no longer crashes in the aforementioned scenario.

BZ#953165


Every time when a user account was saved, the SSSD daemon performed an unnecessary search using an attribute, which was not indexed. As a consequence, saving a large number of user accounts consumed almost 100% of CPU especially during enumeration, because SSSD was searching for a non-indexed attribute. With this update, the search has been disabled and SSSD no longer consumes that amount of CPU when saving enumeration results.

BZ#954275

When an attempt to locate servers using Domain Name System (DNS) Service records (SRV) failed, the SSSD daemon did not retry the SRV query, even when the query internal timeout had passed. Consequently, when the server discovery process failed for the first time, especially during boot up, SSSD did not retry its query until it was restarted or the networking status of the client changed resetting the SSSD networking status. This update applies a patch to fix this bug so that the SRV queries always retry after a timeout passed. As a result, SSSD now retries SRV queries correctly in the described scenario.

BZ#954323

The grace warning code displays number of logins left before the forced change of a password. Previously, there was an “off-by-one” comparison bug in that code. As a consequence, when the 389 Directory Server was used as a server, the last grace warning was not displayed. With this update, the comparison has been fixed and all logins during the grace period now produce warnings as expected.

**NOTE**

Note that the grace warnings currently work only when 389 Directory Server or Red Hat Directory Servers is a Lightweight Directory Access Protocol (LDAP) server. The grace warnings do not work with an OpenLDAP server; this is a known issue.

BZ#963235

When a group whose members were all outside the configured search bases was searched, the search request terminated incorrectly. This caused a use-after-free memory access and therefore the `sssd_be` process could terminate unexpectedly. The search request has been fixed so that it now terminates correctly, even if all group members are outside the configured search bases. As a result, `sssd_be` no longer crashes in the described scenario.

BZ#966757

The default Domain Name System (DNS) timeout values were set too high preventing the SSSD daemon from failing over to all configured DNS servers. When a faulty DNS server was configured in the `/etc/resolv.conf` file, the DNS request could be terminated before it was able to perform a failover through all DNS servers configured in the file. The default DNS timeouts have been lowered allowing SSSD to fail over through all configured DNS servers as expected.

BZ#967636

The number of the autofs maps returned from the SSSD daemon to the automounter daemon was incorrect under certain circumstances, for example, when the maps were too large. As a consequence, the maps were not returned reliably to automounter. This bug has been fixed with this update and the number of maps is now correct in all cases.

BZ#973345

In case the cache contained two entries with the same name, which is an unexpected condition, the search request was not terminated correctly. In fact, the request was terminated twice and therefore the second time the request was terminated, it could access random memory. The error handling during the cache search has been amended so that the request is terminated only once. As a result, the SSSD daemon is now able to handle situations when the cache is corrupted.

BZ#978966

Previously, the sudo refresh handler used an incorrect callback. As a consequence, an incorrect memory could be accessed in certain cases and therefore the `sssd_sudo` process terminated unexpectedly. With this update the handler uses the correct callback so that the process no longer crashes in the described scenario.

BZ#978994

Previously, the description of the “`min_id`” option in the `sssd.conf(5)` manual page was misleading. It stated that the option could be set to the “0” value, which was not correct. With this update, the description has been changed so that the manual page now properly describes that the minimum value for the option is “1”.

BZ#979046

Previously, the IPA provider attempted to store the original value of a member attribute to the cache during the Host-Based Access Control (HBAC) evaluation. The values were processed by the memberof plug-in, which required a lot of processing time when there were very large host groups. As a consequence, the `sssd_be` process used 99% of CPU, which slowed down the login process significantly. With this update, the member attribute is no longer stored and the HBAC evaluation proceeds faster.

BZ#983028

When users attempted to change their password using the `passwd` utility and wrote the current password incorrectly, the following `passwd` error was returned:

```
Authentication token manipulation error
```

This message appeared to be a system error, which could confuse users. With this update, SSSD sends an additional error message that specifies the problem:

```
Old password not accepted
```

BZ#984814

Under certain circumstances, records stored in the fast in-memory cache could become corrupt. In such a case, the `sssd_nss` process terminated unexpectedly. An additional test has been added to check the fast cache before accessing a request. Now, when the records are invalid, they are skipped and requested from the SSSD daemon, thus avoiding the crash of `sssd_nss`.

BZ#986379

Previously, the `sss_cache -N` command did not invalidate the SSSD in-memory cache of netgroups. Consequently, netgroups that had been recently queried were not refreshed before their expiration time, even if the command was executed. This update applies a patch to fix this bug so that the command now correctly invalidates the netgroups in-memory cache.

BZ#987479

The `libsss_sudo` package did not require the certain version of the `sudo` utility that was supposed to work with the SSSD daemon. As a consequence, the package could be installed with the `sudo` version that was not compatible with SSSD. With this update, the package now requires the proper version of `sudo` as expected.

BZ#988525

In case the SSSD daemon could not save a `sudo` rule to the cache, it returned an error and stopped processing the rest of the `sudo` rules. Therefore, none of the rules from the related provider were saved because the error with one rule canceled the entire transaction. With this update, when a `sudo` rule cannot be saved to the cache, a message is appended to the logs and the rule is skipped and processing of the remaining rules continues and works as expected. As a result, all but the defective `sudo` rule are saved to the cache.

BZ#997406

Due to a bug in the underlying source code, a pointer to entries could be overwritten under certain circumstances. Consequently, the `sssd_nss` process terminated unexpectedly with a segmentation fault. The code has been modified to fix this bug and `sssd_nss` no longer crashes.

BZ#1002161

When a large amount of sudo rules with a combined size that exceeded 265 KB was configured on the system, due to the way the `sss_packet_grow()` function computed the total length of a response packet, the SSSD daemon failed with the following error message:

```
Unable to create response: Invalid argument
```

With this update, the `sss_package_grow()` function code has been fixed to properly compute the response packet length, and SSSD no longer fails in the aforementioned scenario.

BZ#1002929

When a dynamic Domain Name System (DNS) update operation timed out, certain data related to the operation was freed. Then a child handler attempted to access those data, which caused a segmentation fault in the `sssd_be` process. This update applies a patch to fix this bug and the handler is now aborted when the operation timed out. As a result, the segmentation fault no longer occurs in the described scenario.

BZ#1019979

In case that the Lightweight Directory Access Protocol (LDAP) connection was terminated when the search operation on this connection was still in progress, the search callback could access properties of the connection that no longer existed. As a consequence, the `sssd_be` process terminated unexpectedly. To fix this bug, an additional test has been added to the search callback. The test checks the validity of a connection before accessing its properties. As a result, the SSSD daemon no longer crashes in the described scenario.

Enhancements

BZ#921454

This update provides a new SSSD configuration option. When enabled, the option permits LDAP groups to contain local users stored in the `/etc/passwd` file. The option is disabled by default, to enable it, set `"ldap_rfc2307_fallback_to_local_users = True"`.

BZ#970519

A new option, which is used to avoid downloading group members, has been introduced. In most cases, the administrator only needs to retrieve group memberships for the user, not to download all group members. Moreover, when the group members are not downloaded and stored to the cache, the SSSD performance increases significantly. With this enhancement, the administrator can now disable downloading group members.

Users of `sssd` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.209. SUBSCRIPTION-MANAGER

8.209.1. [RHBA-2013:1659](#) — subscription-manager and python-rhsm bug fix and enhancement update

Updated `subscription-manager` and `python-rhsm` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The python-rhsm packages provide a library for communicating with the representational state transfer (REST) interface of Red Hat's subscription and content service. The Subscription Management tools use this interface to manage system entitlements, certificates, and content access.

The subscription-manager packages provide programs and libraries to allow users to manage subscriptions and yum repositories from the Red Hat Entitlement platform.



NOTE

- The python-rhsm packages have been upgraded to upstream version 1.9.6, which provides a number of bug fixes and enhancements over the previous version. (BZ#[922837](#))
- The subscription-manager packages have been upgraded to upstream version 1.9.11, which provides a number of bug fixes and enhancements over the previous version. (BZ#[950118](#))
- The subscription-manager-migration-data packages have been upgraded to upstream version 2.0.5, which provides a number of bug fixes and enhancements over the previous version. (BZ#[950116](#))

Bug Fixes

BZ#[1000145](#)

Previously, the python-rhsm utility used a deprecated API. Consequently, a deprecation warning message was displayed to the user. With this update, the deprecation warning message is no longer displayed.

BZ#[914113](#)

Prior to this update, the rhsmc daemon called the deprecated "hasNow()" function. As a consequence, the "DeprecationWarning: Call to deprecated function: hasNow" warning was displayed to the user. With this update, the "hasNow()" function has been removed and the deprecation warning message is no longer displayed.

BZ#[1012566](#)

Prior to this update, the script for the /etc/cron.daily/rhsmc cron job had incorrect permissions. Consequently, even non-root users had execute permissions. This update changes the permissions to the correct "0700" value and only the root user now has execute permissions.

BZ#[872697](#)

Previously, the Japanese translation of the "Configure Pro_xy" message contained an excessive underscore character. Consequently, an incorrect text was displayed to the users of ja_JP locale. This update adds the correct message.

BZ#[985090](#)

Prior to this update, automatic completion of the "rhsmcrtcd" command by pressing the "TAB" key twice did not work properly. Consequently, incorrect options were displayed. The tab completion script has been fixed to display correct options. Note that the bash-completion auxiliary package is required for the auto-completion functionality.

BZ#[988085](#)

Previously, after running the "subscription-manager config --remove <server.hostname>" command,

the "hostname =" line was completely removed from the "rhsm.conf" configuration file. Consequently, the default value of "subscription.rhn.redhat.com" became inaccessible from the command-line interface (CLI). With this update, the "hostname =" line reverts to the expected default value in the described scenario.

BZ#996993, BZ#1008557

This update adds two new fields to the output of the "subscription-manager list --available" command. The "Provides" field shows the names of the products that the system is eligible for. The "Suggested" field has been added to facilitate compliance and provide parity with the graphical user interface (GUI).

BZ#869046

Previously, the subscription-manager utility contained only general error messages when a connection to a proxy failed. As a consequence, users received an uninformative error message when they tried to access an incorrect proxy server, tried to connect via an incorrect proxy port, or failed to enter the correct password. This update adds more informative error messages for the described cases.

BZ#1001820

Prior to this update, automatic completion of the "subscription-manager attach" subcommand by pressing the "TAB" key twice did not work properly. As a consequence, incorrect options were displayed. The tab completion script has been fixed to display correct options. Note that the bash-completion auxiliary package is required for the auto-completion functionality.

BZ#1004385

Previously, automatic completion of the "rhsm-icon" command by pressing the "TAB" key twice did not work properly. Consequently, options were displayed with a comma at the end. The tab completion script has been fixed to display correct options. Note that the bash-completion auxiliary package is required for the auto-completion functionality.

BZ#1004893

Under certain circumstances, the "subscription-manager list --installed" command returned an incorrect status. Consequently, when a new product certificate contained a new product, the displayed status of the newly available product was "Not Subscribed". This bug has been fixed and the displayed status for the newly available product is now "Subscribed" in the described scenario.

BZ#1011234

Under certain circumstances, the "subscription-manager list --available" command returned an incorrect value. Consequently, for subscription pools whose Service Level had not been set, misleading "None" was displayed. This bug has been fixed and an empty string is now displayed in this scenario.

BZ#1006985

Prior to this update, the subscription-manager-migration script did not work properly when migrating different product certificates with the same product ID. As a consequence, the certificates were installed under the same name and were unusable. This bug has been fixed and the migration is aborted when different product certificates with the same ID are detected.

BZ#1008603

Previously, the subscription-manager utility required connectivity to the "subscription.rhn.stage.redhat.com" site in order to list products. Consequently, the product list was

not displayed when the connection failed. This bug has been fixed and users are now able to list products from the local cache.

Enhancements

BZ#909778

This update adds the "--proxy" option to the "subscription-manager repos --list" subcommand. The user is now able to set the proxy when connecting to the candlepin server.

BZ#983670

The description displayed when using the "--help" option with the "subscription-manager auto-attach" subcommand has been improved to be more precise.

BZ#986971

The "Available Subscriptions" header in the Subscriptions table has been simplified to just "Available", which saves space and is clearer to the user.

BZ#1011961

With this update, the displayed quantity in the Entitlement Certificate has been changed from the confusing "-1" to the correct "Unlimited".

BZ#994620

This update provides a more precise tooltip messaging for the rhsm-icon utility. Now, when a partial subscription exists on a fully compliant machine, the message says "Partially entitled subscriptions" instead of the previous "Partially entitled products".

BZ#1004341

This update adds support for automatic completion of the "subscription-manager-gui" command options by pressing the "TAB" key twice. Note that the bash-completion auxiliary package is required for the auto-completion functionality.

BZ#1008016

With this update, the subscription-manager utility generates the /etc/yum.repos.d/redhat.repo repository immediately after a successful subscription, no more steps are necessary.

BZ#1009600

When the "subscription-manager list --consumed" command is run, the output now displays "System Type: Physical/Virtual". This allows the user to determine whether the granted entitlement was virtual.

Users of subscription-manager and python-rhsm are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.210. SUDO

8.210.1. RHSA-2013:1701 — Low: sudo security, bug fix and enhancement update

An updated sudo package that fixes two security issues, several bugs, and adds two enhancements is now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The sudo (superuser do) utility allows system administrators to give certain users the ability to run commands as root.

Security Fixes

[CVE-2013-1775](#)

A flaw was found in the way sudo handled time stamp files. An attacker able to run code as a local user and with the ability to control the system clock could possibly gain additional privileges by running commands that the victim user was allowed to run via sudo, without knowing the victim's password.

[CVE-2013-2776](#), [CVE-2013-2777](#)

It was found that sudo did not properly validate the controlling terminal device when the `tty_tickets` option was enabled in the `/etc/sudoers` file. An attacker able to run code as a local user could possibly gain additional privileges by running commands that the victim user was allowed to run via sudo, without knowing the victim's password.

Bug Fixes

[BZ#880150](#)

Previously, sudo did not support netgroup filtering for sources from the System Security Services Daemon (SSSD). Consequently, SSSD rules were applied to all users even when they did not belong to the specified netgroup. With this update, netgroup filtering for SSSD sources has been implemented. As a result, rules with a netgroup specification are applied only to users that are part of the netgroup.

[BZ#947276](#)

When the sudo utility set up the environment in which it ran a command, it reset the value of the `RLIMIT_NPROC` resource limit to the parent's value of this limit if both the soft (current) and hard (maximum) values of `RLIMIT_NPROC` were not limited. An upstream patch has been provided to address this bug and `RLIMIT_NPROC` can now be set to "unlimited".

[BZ#973228](#)

Due to the refactoring of the sudo code by upstream, the `SUDO_USER` variable that stores the name of the user running the sudo command was not logged to the `/var/log/secure` file as before. Consequently, user name "root" was always recorded instead of the real user name. With this update, the previous behavior of sudo has been restored. As a result, the expected user name is now written to `/var/log/secure`.

[BZ#994626](#)

Due to an error in a loop condition in sudo's rule listing code, a buffer overflow could have occurred in certain cases. This condition has been fixed and the buffer overflow no longer occurs.

Enhancements

[BZ#848111](#)

With this update, sudo has been modified to send debug messages about netgroup matching to the debug log. These messages should provide better understanding of how sudo matches netgroup database records with values from the running system and what the values are exactly.

BZ#853542

With this update, sudo has been modified to accept the `ipa_hostname` value from the `/etc/sss/sss.conf` configuration file when matching netgroups.

All sudo users are advised to upgrade to this updated package, which contains backported patches to correct these issues and add these enhancements.

8.211. SUITESPARSE

8.211.1. RHBA-2013:0988 — suitesparse bug fix update

Updated suitesparse packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The suitesparse packages are a collection of libraries for computations involving sparse matrices.

Bug Fix

BZ#902854

The suitesparse base package was missing a copy of the General Public License (GPL). The license was included in the suitesparse-doc subpackage, but it was possible to install suitesparse without the suitesparse-doc subpackage. With this update, a copy of the license is now included also in the base package.

Users of suitesparse are advised to upgrade to these updated packages, which fix this bug.

8.212. SYSSTAT

8.212.1. RHBA-2013:1663 — sysstat bug fix and enhancement update

Updated sysstat packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The sysstat packages provide a set of utilities which enable system monitoring of disks, network, and other I/O activity.

Bug Fixes

BZ#804534

Previously, the sysstat package did not support dynamically attributed major device numbers. Consequently, devices with these numbers were not listed in sar reports under their real names. With this update, support for dynamically attributed major device numbers has been added to sysstat. As a result, all devices now appear with their correct names in sar reports.

BZ#967386

A previous sysstat update changed binary data files in a backward incompatible way, but the version number of these binary data files remained the same. Consequently, using a later sysstat version to

read binary data files created by an earlier version of sysstat could have produced invalid results. The version number of sysstat binary data files has been updated, thus fixing this bug. As a result, the current sysstat version will not read binary data files created by previous versions. For more information, please refer to the description of the "--legacy" option in the sar(1) manual page.

BZ#996134

Prior to this update, the umask command was executed too late in the sa1 script. Under certain circumstances, this could have caused incorrect file permissions of newly created files. With this update, executing umask has been moved to the appropriate place in the sa1 script. As a result, newly created files have correct permissions.

Enhancements**BZ#826399**

Kernel device names, such as sda or sdb, might point at different devices every boot. To prevent possible confusion, support for persistent device names has been added to the iostat and sar programs. Persistent names can be enabled with the new "-j" command-line option for both iostat and sar.

BZ#838914

The sysstat package has been modified to store the collected statistics longer. The original period of 7 days has been extended to 28 days, thus allowing for better analysis of more complex performance issues.

BZ#850810

With this update, a new "-y" option has been added to the iostat program. This option allows to skip first "since boot" statistics in the report, so there is no longer need to post-process the iostat output in this matter.

Users of sysstat are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.213. SYSTEM-CONFIG-DATE

8.213.1. [RHBA-2013:1098](#) — [system-config-date bug fix update](#)

An updated system-config-date package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The system-config-date package provides a graphical interface for changing the system date and time, configuring the system time zone, and setting up the NTP daemon to synchronize the time of the system with an NTP time server.

Bug Fix**BZ#760977**

When using unsupported locale settings during system installation, the firstboot utility previously failed with a stack trace after the initial phase of installation. Consequently, the user was not able to configure peripherals, users, kdump, and other settings. With this update, the underlying code has

been modified to catch exceptions caused by incorrect locale values and firstboot now warns the users to change their locale instead of failing.

Users of system-config-date are advised to upgrade to this updated package, which fixes this bug.

8.214. SYSTEM-CONFIG-KEYBOARD

8.214.1. RHBA-2013:0940 — system-config-keyboard bug fix update

Updated system-config-keyboard packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The system-config-keyboard packages provide a graphical user interface that allows the user to change the default keyboard of the system.

Bug Fix

BZ#952125

The system-config-keyboard packages contain a plug-in for firstboot. Previous versions of system-config-keyboard depended on firstboot, so it was not possible to install the packages without pulling in firstboot too. This erroneous dependency has been removed and the system-config-keyboard packages can now be installed without pulling in firstboot.

Users of system-config-keyboard are advised to upgrade to these updated packages, which fix this bug.

8.215. SYSTEM-CONFIG-LVM

8.215.1. RHBA-2013:1570 — system-config-lvm bug fix update

An updated system-config-lvm package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The system-config-lvm utility enables users to configure logical volumes using a GUI.

Bug Fix

BZ#923643

Due to a bug in the system-config-lvm utility, the utility terminated unexpectedly when striped mirrored devices were found on the system. With this update, the underlying source code has been modified so that the users can now fully interact with supported devices. However, volume group information may not always render properly for striped mirrored devices.

Users of system-config-lvm are advised to upgrade to this updated package, which fixes this bug.

8.216. SYSTEM-CONFIG-USERS-DOCS

8.216.1. RHBA-2013:1002 — system-config-users-docs bug fix update

An updated system-config-users-docs package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The system-config-users-docs package contains the online documentation for system-config-users, which is a graphical utility for administrating users and groups.

Bug Fix

BZ#635248

Previously, text in screenshots was not translated and the untranslated screenshots did not match the running program. This update brings the screenshots up to date and their translations are now included.

Users of system-config-users-docs are advised to upgrade to this updated package, which fixes this bug.

8.217. SYSTEMTAP

8.217.1. [RHBA-2013:1630](#) — systemtap bug fix and enhancement update

Updated systemtap packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

SystemTap is a tracing and probing tool to analyze and monitor activities of the operating system, including the kernel. It provides a wide range of filtering and analysis options.



NOTE

The systemtap packages have been upgraded to upstream version 2.3, which provides a number of bug fixes and enhancements over the previous version.

The most notable enhancements are:

- * new support for regular expression matching in script language, using the "=~" and "!~" operators;
- * improved error diagnostics with references to man pages that give further information and corrective advice;
- * new ability to read the perf-counter values in process probes, using the "@perf()" operator;
- * faster stack unwinding functions, especially for partial "unwinds" of only one or a few call levels;
- * new support for a macro processing in the scripting language, using the "@define" operator;
- * various tapset and runtime improvements.

(BZ#[920682](#))

Bug Fixes

BZ#[920444](#)

Previously, the SystemTap "nfs.proc.commit_done" probe alias interpreted the underlying arguments incorrectly. As a consequence, incorrect information was returned. The SystemTap "nfs_proc.commit_done" probe alias has been modified to interpret the underlying arguments correctly, and it now returns the correct information.

BZ#743591

Prior to this update, the ioblktime.stp example script did not track properly when I/O block pending requests were merged together with other requests. Consequently, the associative array could overflow, and an error occurred. The script has been modified to delete the merged I/O block requests. As a result, the script no longer fails after array overflow errors.

BZ#874205

The stap-serverd daemon did not use the avahi-client API properly. As a consequence, stap-serverd terminated unexpectedly when the avahi daemon was stopped or restarted. The underlying source code has been modified to use the avahi-client API properly, and stap-serverd now handles the described scenario without crashing.

BZ#876848

If the main thread of an application did not go through at least one quiesce, such as a system call, the SystemTap utility was unable to attach to the application. Consequently, SystemTap could not detect the events of the application. SystemTap now interrupts the main application thread so that it can attach to it. As a result, Systemtap no longer misses any events.

BZ#846789

Previously, the SystemTap utility used an arbitrary order when searching for kernel modules under the /lib/modules/ directory and identified the base module first. As a consequence, SystemTap could not probe some modules that were overridden by updates. SystemTap now uses the modules.dep file to order the module searches so that they match the order of the modprobe program. As a result, the overridden modules can be probed successfully.

BZ#819967

Previously, the SystemTap translator did not always include the correct code to be able to use the caller() function, which could lead to compile errors. This update ensures that the SystemTap translator includes the proper runtime code when using the caller() function. As a result, using caller() no longer causes compile errors.

BZ#906061

Prior to this update, the scope specifiers and IDs for IPv6 addresses were not handled correctly in SystemTap. Consequently, the user could not connect to stap-server using IPv6. This update corrects the handling of scope specifiers or IDs for IPv6 addresses, and the user can now connect to stap-server using IPv6 as expected.

BZ#902739

Previously, the "trace" utility in SystemTap generated a low-quality C code. As a consequence, when the generated code was compiled, the gcc compiler displayed error messages. This update improves the quality of the generated code, and no warnings are now displayed as expected.

Users of systemtap are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.218. SYSVINIT

8.218.1. RHBA-2013:1003 — sysvinit bug fix update

Updated sysvinit packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The sysvinit packages contain a group of processes that control basic functions of your system. SysVinit includes the init program, which is the first program started by the Linux kernel when the system boots. The init program then controls the starting up and shutting down of all other programs.

Bug Fix

BZ#814132

When the pidof utility was processing the /proc/*/stat file, the content of the file was read by the fgets() function. This function did not behave correctly when the new line character, "\n", was in the name of a process and also when the read file disappeared between the fopen and fgets function. Consequently, the "pidof: could not get program name from" message was emitted in these cases. To fix this problem, the program reads the whole file and silently skips all files with empty content. As a result, binary files containing the new line character with "\n" are now correctly recognized and disappearing files do not emit any warning messages.

Users of sysvinit are advised to upgrade to these updated packages, which fix this bug.

8.219. TALK

8.219.1. RHBA-2013:1148 — talk bug fix update

Updated talk packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The talk utility is a communication program that copies lines from one terminal to the terminal of another user.

Bug Fix

BZ#691355

The talk utility allows a user to specify the target user in the "username.hostname" form. Consequent to this, previous versions of the utility did not support usernames that contained a period. With this update, a new command line option (that is, "-x") has been added to enforce the use of the "username@hostname" form, so that the username can contain periods. As well, the corresponding manual page has been extended to provide a complete list of supported command line arguments.

Users of talk are advised to upgrade to these updated packages, which fix this bug.

8.220. TBOOT

8.220.1. RHBA-2013:1606 — tboot bug fix and enhancement update

Updated tboot packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The tboot packages provide the Trusted Boot (tboot) open source pre-kernel/VMM module. This module uses Intel Trusted Execution Technology (Intel TXT) to initialize the launch of operating system kernels and virtual machines.



NOTE

The tboot packages have been upgraded to upstream version 1.7.4, which provides a number of bug fixes and enhancements over the previous version. (BZ#[916046](#), BZ#[957158](#))

Users of tboot are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.221. TOMCAT6

8.221.1. RHBA-2013:1721 — tomcat6 bug fix update

Updated tomcat6 packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

Apache Tomcat is a servlet container for the Java Servlet and JavaServer Pages (JSP) technologies.

Bug Fixes

BZ#[845786](#)

Previously, an attempt to build the tomcat6-docs-webapp package failed when Red Hat Enterprise Linux was running on IBM System z or 64-bit IBM POWER Series computers. With this update, no architecture is set in the build target and the package can be built as expected.

BZ#[915447](#)

When a user, whose name did not correspond to any existing group name, was specified in the `/etc/sysconfig/tomcat6` file, the Tomcat web server failed to start. This update applies a patch to fix this bug and Tomcat no longer fails in the described scenario.

BZ#[950647](#)

Due to a bug in the `checkpidfile()` function, an attempt to execute the "service tomcat6 status" command failed and an error message was returned. The underlying source code has been modified to fix this bug and the command now works properly.

BZ#[960255](#)

Due to a bug in the `checkpidfile()` function, the status script did not return the correct PID. This bug has been fixed and the status script now returns the correct PID as expected.

BZ#[977685](#)

The Tomcat web server included a version of the `tomcat-juli.jar` file that was hard coded to use classes from the `java.util.logging` package instead of the `log4j` framework. Consequently, Tomcat could not be configured to use `log4j` unless the complete version of the `tomcat-juli.jar` and `tomcat-juli-adapters.jar` files had been downloaded. With this update, the tomcat6 packages now contain the correct versions of these files to configure `log4j`.

BZ#[989527](#)

When multiple tomcat instances were configured as described in the `/etc/sysconfig/tomcat6` configuration file and the instance name was different from the name of the tomcat directory, the "service status" command failed. With this update, the underlying source code has been modified to fix this bug and the command no longer fails in the described scenario.

Users of tomcat6 are advised to upgrade to these updated packages, which fix these bugs.

8.222. TUNED

8.222.1. [RHBA-2013:1623](#) — tuned bug fix and enhancement update

Updated tuned packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The tuned packages contain a daemon that tunes system settings dynamically. It does so by monitoring the usage of several system components periodically.

Bug Fixes

[BZ#904062](#)

Previously, when multiple devices were added into the system, a udev rule restarted the ktune service for each new device. This could lead to many restarts in a short period of time. The multiple restarts could trigger a race condition in the kernel, which cannot be currently fixed. The tuned daemon code has been modified not to trigger more than one restart per 10 seconds, thus preventing the race condition from occurring.

[BZ#969491](#)

The `kernel.sched_migration_cost` tunable was previously kept at its default value of 0.5 ms. As a consequence, big virtualization hosts could experience high contention at the run queue lock. With this update, the `kernel.sched_migration_cost` tunable has been increased 10 times in the virtual-host profile, which eliminates the contention.

[BZ#905077](#)

Previously, the ktune service did not save readahead values. On startup, it multiplied the current value by a constant and divided the value by the same constant on stop. This could result in a wrong value being set on devices that were added after ktune had been started. Now, the previous readahead values are stored for all devices and the correct values are restored on ktune stop.

[BZ#912788](#)

Previously, the tuned utility did not support the upstream `/sys/kernel/mm/transparent_hugepage` location for Transparent Huge Pages (THP). The code has been modified and support for the `SYSFS` path used by upstream has been added. Tuned now supports the aforementioned upstream path as well as the Red Hat Enterprise Linux specific one.

[BZ#982756](#)

There was a typo in the USB autosuspend code and also an old non-functional code in the function that handles Bluetooth. As a consequence, various errors occurred when the user activated the spindown-disk profile. The typo has been fixed and the Bluetooth code has been updated. As a result, errors no longer occur when the spindown-disk profile is activated.

[BZ#838512](#)

Previously, the mount command was used to remount file systems with the "no_barriers" option, but not all file systems could be remounted. As a consequence, the remount occasionally failed displaying an error message, which could confuse users. With this update, the error messages from the mount command have been silenced. Now, if the file system cannot be remounted with no_barriers, it is silently skipped and no error is displayed.

BZ#987547

Previously, the sysctl utility was loaded and the shell script was run before the elevator was changed. As a consequence, the user was unable to change or tune the elevator parameters. The code has been reordered to load sysctl and run the shell script after the elevator has been changed. As a result, the user can now tune the elevator parameters.

BZ#885080

The diskdevstat and netdevstat code was not consistent in naming the parameters in the built-in help. The terms "total-duration" and "total-interval" have been used to denote the same thing. With this update, the text has been updated to be consistent and now, only the "total-duration" string is used.

BZ#959732

Previously ktune did not handle the /etc/sysctl.d/ directory and did not support multiple files with sysctl settings. As a consequence, the settings that several packages, for example libvirt, installed in the /etc/sysctl.d directory, were ignored. After this update, the ktune code and profiles have been modified. Now all sysctl settings in the /etc/sysctl.d/ directory are loaded and then the /etc/sysctl.conf file is loaded. The user can now specify multiple sysctl files, including wildcards, to load in the tuned profiles.

BZ#964187

Previously, there was no documentation for the Tuned virtual-guest and virtual-host profiles. Descriptions for these profiles have been added to the tuned-adm manual page.

BZ#963821

Previously, there was a typo in the built-in help of the tuned-adm command. The text mentioned "tunning" instead of "tuning", in the description of tuned-adm. This update fixes the typo in the built-in help.

BZ#961792

Tuned previously locked the CPU to the C0 state in the latency-performance profile. With this update, the latency-performance profile has been modified to use the C1 state by default, thus improving the performance of the profile.

Enhancements**BZ#964193**

This update adds the "sapconf" package to the Tuned utility. The tuned-adm "sap" profile is used to optimize systems running SAP software, matching the relevant SAP guidelines.

Users of tuned are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.223. UDEV

8.223.1. RHBA-2013:1675 — udev bug fix and enhancement update

Updated udev packages that fix several bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The udev packages implement a dynamic device-directory, providing only the devices present on the system. This dynamic directory runs in user space, dynamically creates and removes devices, provides consistent naming, and a user-space API. The udev packages replace the devfs package and provides better hot plug functionality.

Bug Fixes

BZ#833172, BZ#885978, BZ#918511

Previously, for machines with relatively big RAM sizes and lots of disks, a number of udevd workers were running in parallel, maximizing CPU and I/O. This could cause udev events to time out due to hardware bottlenecks. With this update, the number of udevd workers is limited by the CPU count and significantly lower on machines with a big RAM size. Now, fewer udev workers running concurrently do not bottleneck easily and cause less or no timeouts.

BZ#888647

Previously, the udev utility did not provide a symbolic link to SCM (Storage Class Memory) devices in the `/dev/disk/by-path/` directory, which prevented SCM devices to be referenced by their paths. With this update, the `path_id` built-in command supports SCM devices and provides a symbolic link. Now, SCM devices can be referenced by their paths.

BZ#909792

Prior to this update, the `libudev.h` header file did not have any extern "C" declaration, so it could not be used as-is in a C++ programs or applications. An extern "C" declaration has been added to the header file, thus fixing the bug.

BZ#918511

Previously, the `start_udev` command called the "udevadm settle [options]" command and timed out after the default of 180 seconds. Nevertheless, some devices were not completely assembled and the boot process continued causing various failures. With this update, `start_udev` waits until udev has settled. As a result, all devices are assembled, and the boot process now continues without errors.

BZ#920961

If a SCSI device was in use at the time the udev `scsi_id` helper utility was invoked, `scsi_id` did not return any properties of the device. Consequently, the properties of the SCSI device could not be processed in udev rules. With this update, `scsi_id` retries to open the device for a certain time span before it gives up. As a result, the properties of a SCSI device can be processed in udev rules, even though the device is in use for a short time.

BZ#982902

For USB devices with `InterfaceClass=0x08` and `InterfaceSubClass=0x05`, udev set the ID type as "floppy", which was not necessarily true. As a consequence, some tools could interpret the USB device as a floppy disk. Now, the ID type is set as "generic" for such USB devices, and tools interpret the USB devices correctly.

BZ#998237

Previously, the `libudev` utility was referencing memory, which had been reallocated with its old address into the `dev_enumerate_get_list_entry()` function. However, calling this function could lead to

a segmentation fault. With this update, libudev references the reallocated memory with offsets in `udev_enumerate_get_list_entry()`, thus fixing the bug.

Enhancement

BZ#947067

Previously, the amount of debug output could not be controlled and often exceeded the available memory, if stored in the `/dev/` temporary file. With this update, the `start_udev` command with `udevlog` now call the `udev` daemon with the `-s` option, which redirects the output of `udev` to the `/dev/.udev/udev.log` file but does not set `udev` in the debug mode. In addition, `udev` now understands the log priorities set in the rules file (`OPTIONS+="log_priority=<level>"`), so the user can set the numerical syslog priorities or their textual representations. There is also a new example rules file for logging: `/lib/udev/rules.d/01-log-block.rules`. To enable "info" logging for block devices, add `"rd.log.block=info"` to the kernel command line.

Users of `udev` are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.224. UTIL-LINUX-NG

8.224.1. RHBA-2013:1648 — util-linux-ng bug fix and enhancement update

Updated `util-linux-ng` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `util-linux-ng` packages contain a large variety of low-level system utilities that are necessary for a Linux operating system to function.

Bug Fixes

BZ#885313

Previously, the `hexdump` utility terminated with a segmentation fault when iterating over an empty format string. This bug has now been fixed and `hexdump` no longer crashes in this scenario.

BZ#911756

Previously, the `libblkid` library incorrectly detected certain disks as a Silicon Image Medley RAID device. Consequently, this caused problems in certain systems after a weekly reboot. This update adds a checksum count from the superblock record and a new superblock definition from the `dmraid` tool, which makes the signature recognition of Silicon Image Medley RAID devices more robust.

BZ#864585

Previously, the `"mount -av"` command, which triggers mounting filesystems with helpers like the `/sbin/mount.nfs` file, printed the message "nothing was mounted", even though the helper mounted a filesystem. This bug has been fixed and the incorrect message is no longer printed in this scenario.

BZ#872291

Previously, the `hwclock(8)` manual page contained a reference to the non-existing `adjtimex` utility. This update fixes the `hwclock(8)` manual page.

BZ#915844

Previously, the `mount(8)` manual page incorrectly described the "relatime" mount option. With this update, the description of the "relatime" mount option has been improved to better describe when the kernel updates the atime.

BZ#917678

Due to a regression in the code, if a symbolic link was used for a mount point in the `/etc/fstab` configuration file, mount attempts to that mount point failed. This update ensures that all paths in `/etc/fstab` are made canonical and such mount points can now be mounted as expected.

BZ#966735

Prior to this update, the `lscpu` command accepted only sequentially assigned logical CPU numbers. Consequently, `lscpu` did not properly list CPUs after a CPU eject operation. After this update, the `lscpu` command does not expect sequentially assigned CPU numbers and works properly on systems with a hot-plug CPU.

Enhancements**BZ#816342**

Previously, it was not possible to determine the right CLOCAL flag by the kernel and, also, some machines required manual settings. With this update, the new `-L[={always,auto,never}]` option has been added to the `agetty` utility to allow complete control on the CLOCAL terminal flag.

BZ#846790

Previously, the `kill(1)` manual page did not include information about the interaction between the `kill` utility and threads. With this update, the `kill(1)` manual page has been improved to explicitly explain the interaction between the `kill` system call and threads.

BZ#870854

The default kill character "@" was in collision with login user names on IPA systems with the "user@domain" convention. With this update, the `agetty` utility has been improved to accept the "--kill-chars" and "--erase-chars" options to control special kill and erase terminal characters.

BZ#947062

With this update, the "blkdiscard" command has been introduced to Red Hat Enterprise Linux 6 to discard device sectors. The "discard" support is important for example on thinly-provisioned storage to improve disk efficiency by reclaiming free space so that the storage can re-use the free space for other area.

Users of `util-linux-ng` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.225. VHOSTMD

8.225.1. RHBA-2013:1579 — vhostmd bug fix update

Updated `vhostmd` packages that fix one bug are now available for Red Hat Enterprise Linux 6 for SAP.

The Virtual Host Metrics Daemon (`vhostmd`) provides virtual machines with information on the resource utilization of the Red Hat Enterprise Linux host on which they are being run.

Bug Fix

BZ#820500

Due to bugs in the libmetrics code, user's programs could terminate with a segmentation fault when attempting to obtain guest metrics from vhostmd. The libmetrics code has been fixed to perform XPath queries and propagate errors to the user correctly so that the user's programs can now obtain guest metrics as expected.

All users of vhostmd are advised to upgrade to these updated packages, which fix these bugs.

8.226. VIRT-MANAGER

8.226.1. RHBA-2013:1646 — virt-manager bug fix update

Updated virt-manager packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

Virtual Machine Manager (virt-manager) is a graphical tool for administering virtual machines for KVM, Xen, and QEMU. The virt-manager utility can start, stop, add or remove virtualized devices, connect to a graphical or serial console, and see resource usage statistics for existing virtualized guests on local or remote machines. It uses the libvirt API (Application Programming Interface).

Bug Fixes

BZ#820303

Previously, when calling the libvirt utility, virt-manager omitted an address (in form "bus:device") when identical USB devices (in form "vendorid:productid") were attached, and thus the wrong devices were attached to the guest. With this update, the user specifies information about both, the "bus:device" and "vendorid:productid", to select the correct device. Now, the specified device in the XML or the device selected in the virt-manager GUI are correctly attached to the guest.

BZ#869206

Previously, changing a device type or model did not reset the guest address that the device should be reachable at. Consequently, the guest could not start after changing a watchdog from i6300esb to ib700. This bug has been fixed and the guest can now be started as expected.

BZ#869474

When selecting a bridge network created by the libvirt utility, virt-manager could not display the details and configuration of network created by libvirt. Moreover, the following error was returned:

```
Error selecting network: 'None Type' object has no attribute 'split'
```

With this update, configuration of the network created by libvirt.

BZ#873142

Previously, the "create a new virtual machine" virt-manager dialog contained a typographical mistake in unit of "Storage", showing "Gb" instead of "GB". The typo has been fixed.

BZ#907399

Due to a wrong attribute always set to "no", errors occurred after changing SELinux from the static option to dynamic on virt-manager. A patch has been provided to fix this bug. With this update, no error messages are returned and SELinux now changes from the static to dynamic option

successfully.

BZ#981628

If the "Toolbar" check-box was unchecked from the VM configuration in virt-manager, any new VM failed to start installation and the 'Begin Installation' button disappeared. A patch has been applied to fix this bug, and the 'Begin Installation' button no longer disappears from the GUI.

BZ#985184

Previously, the ram attribute supported only the qxl guest driver type. Consequently, errors were shown when changing a video from qxl to other models. With this update, the guest works well and deletes the "ram" element automatically when models are changed.

BZ#990507

Prior to this update, using virt-manager to connect a physical CD-ROM or an ISO CD-ROM image occasionally did not work in KDE. Also, the "Choose Media" dialog box to select the image or physical device did not show up. A patch has been provided to fix this bug and the "Choose Media" dialog window now shows up when the "Connect" button is pressed.

All virt-manager users are advised to upgrade to these updated packages, which fix these bugs.

8.227. VIRT-P2V

8.227.1. [RHEA-2013:1631](#) — virt-p2v bug fix and enhancement update

Updated virt-p2v packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Virt-p2v is a tool for conversion of a physical server to a virtual guest.

**NOTE**

The virt-p2v package has been upgraded to upstream version 0.9.1, which provides a number of bug fixes and enhancements over the previous version. ([BZ#976832](#))

Users of virt-p2v are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.228. VIRT-V2V

8.228.1. [RHBA-2013:1547](#) — virt-v2v bug fix and enhancement update

Updated virt-v2v package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The virt-v2v package provides a tool for converting virtual machines to use the KVM (Kernel-based Virtual Machine) hypervisor or Red Hat Enterprise Virtualization. The tool modifies both the virtual machine image and its associated libvirt metadata. Also, virt-v2v can configure a guest to use VirtIO drivers if possible.

**NOTE**

The virt-v2v package has been upgraded to upstream version 0.9.1, which provides a number of bug fixes and enhancements over the previous version. (BZ#965501)

Bug Fixes**BZ#809273**

After using the virt-v2v utility to migrate a Windows 2003 R2 32bit guest from Red Hat Enterprise Linux 5 Xen to Red Hat Enterprise Linux 6 KVM, the guest did not boot. With this update, the Windows Xen drivers are disabled during conversion, and the guest now boots correctly.

BZ#820928

Previously, virt-v2v assumed that Microsoft Windows operating system could be installed only to the /windows directory. Consequently, when a non-existent path was copied during installation, virt-v2v terminated unexpectedly. With this update, the system path is no longer hardcoded, and virt-v2v no longer crashes in the described scenario.

BZ#829859

Previously, virt-v2v always used the first kernel (0th kernel) instead of the default kernel (ex: default=1). With this update, virt-v2v uses the default kernel to create a new migrated Virtual Machine.

BZ#887884

Prior to this update, during conversion of a Linux guest with multiple console entries on the kernel boot command line, virt-v2v terminated unexpectedly. With this update, guests with multiple consoles can be converted successfully.

BZ#953994

Due to the incompatibility of vmware-tools and the virt-v2v code, the virt-v2v utility was failing to convert ESX Red Hat Enterprise Linux VM to KVM VM (Kernel-based VM) on the vmware-tools-foundation dependency. An upstream patch that fixes the broken dependency has been applied, and conversions now succeed.

Users of virt-v2v are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

8.228.2. RHBA-2013:1749 — virt-v2v bug fix update

Updated virt-v2v package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The virt-v2v package provides a tool for converting virtual machines to use the KVM (Kernel-based Virtual Machine) hypervisor or Red Hat Enterprise Virtualization. The tool modifies both the virtual machine image and its associated libvirt metadata. Also, virt-v2v can configure a guest to use VirtIO drivers if possible.

Bug Fix**BZ#1028983**

An update to virt-v2v included upstream support for the import of OVA images exported by VMware servers. Unfortunately, testing has shown that VMDK images created by recent versions of VMware ESX cannot be reliably supported, thus this feature has been withdrawn.

Users of virt-v2v are advised to upgrade to this updated package, which fixes this bug.

8.229. VIRT-VIEWER

8.229.1. RHBA-2013:1578 — virt-viewer bug fix and enhancement update

Updated virt-viewer packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Virtual Machine Viewer (virt-viewer) is a lightweight interface for interacting with the graphical display of a virtualized guest. Virtual Machine Viewer uses libvirt and is intended as a replacement for traditional VNC or SPICE clients.



NOTE

The virt-viewer package has been upgraded to upstream version 0.5.6, which provides a number of bug fixes and enhancements over the previous version, including fixes for issues with multiple monitors. Also, it is now possible to disable a display through the Windows screen resolution utility. (BZ#961455, BZ#888629, BZ#958966)

Bug Fixes

BZ#809546

When a guest was configured with two or more screens, these screens were not shown correctly on the respective client monitors. With this update, virt-viewer has been modified to place the guest screens correctly on the client monitors.

BZ#846127

Due to a bug in the **spice-gtk** utility, when a virtual machine was started in full screen mode, the guest screen resolution could not be changed afterwards. Consequently, the resolution always reverted to the native monitor resolution. This bug has been fixed, and the screen resolution can now be changed without complications.

BZ#856682

Previously, when a Gtk+ menu was open during certain operation that modified its content, an attempt to recreate this menu could cause the client to become unresponsive. With this update, instead of recreating the whole menu, **virt-viewer** repopulates the existing one. As a result, the risk of client freeze is now reduced in the aforementioned scenario.

BZ#864929

Prior to this update, when the client was in full screen mode, it was not possible to set up screen resolution higher than the native resolution of the monitor. The underlying source code has been modified and resolutions higher than the monitor native resolution can now be configured in full screen mode.

BZ#870710

Previously, keyboard events were not sent to the guest when it was suspended, therefore the guest could not be woken up. This bug has been fixed, and all keyboard events are now sent to the guest, regardless of its state. As a result, pressing a key now wakes up the guest as expected.

BZ#875697

When the guest was shut down while the client was still connected, the following message was displayed:

```
Unable to connect to the graphic server
```

With this update, **virt-viewer** has been modified not to report an error on normal disconnection. As a result, the error message is no longer displayed in the described case.

BZ#876444

Previously, when **virt-viewer** ran in full screen mode, the mirror monitors were created randomly. This update modifies **virt-viewer** to maintain the association of client window and monitors. As a result, the additional monitors are no longer mirrored randomly.

BZ#876445

Prior to this update, the title bar of the client window became invisible when **virt-viewer** was leaving the full screen state. This bug has been fixed and the client window title bar is now visible and reachable when leaving full screen.

BZ#886570

Previously, when **virt-viewer** was switched to full screen mode, multiple monitor displays in both client and guest appeared to have the same resolution, even though the monitors were different. This bug has been fixed, and monitors now display resolution according to their actual capacity.

BZ#890297

Prior to this update, after closing the **virt-viewer** guest terminal, the I/O error was written to the **libvirtd.log** file. With this update, **libvirt** events and callbacks are unregistered when closing the guest terminal, and I/O errors are no longer logged in the aforementioned scenario.

BZ#908057

When the **automatically resize** option was disabled in **remote-viewer** and the screen resolution on the guest machine was changed, this change was not accepted and the resolution reverted back to the previous state. With this update, **remote-viewer** has been modified to keep monitor configuration synchronized with the guest, even when automatic resize is disabled.

BZ#908408

Due to an unnecessary message returned by **virt-viewer**, when attempting to connect to a multi-monitor guest in full screen mode, secondary monitors sometimes kept flashing in a loop. Now, the message about additional monitor reconfiguration has been removed and extra monitors no longer flash when in full screen mode.

BZ#913601

Prior to this update, in a multi-monitor guest setup, the sendkey menu was incorrectly placed on the secondary monitors. With this update, **virt-viewer** has been modified to correctly translate the menu coordinates based on the top level window position. As a result, the sendkey menu is now correctly placed on secondary monitors.

BZ#924577

After pressing the **Alt+S** key combination or other menu accelerators the guest kept the **Alt** state enabled. Consequently, certain guest functionality did not work correctly. With this update, the guest **Alt** keys are properly released when the keyboard grab is taken in the client user interface, thus fixing this bug.

BZ#982840

Previously, **virt-viewer** was not able to connect to the libvirt guest console configured with only SPICE TLS autoport. With this update, the setup logic has been modified to check the presence of any port, plain or TLS. As a result, **virt-viewer** can now connect successfully to the TLS-only guest.

BZ#990883

When an invalid password was entered for a Spice session, no error message was shown. Now, the authentication failure is handled properly and an error dialog is shown if Spice password is invalid.

Enhancements**BZ#864026**

This update modifies the **virt-viewer** startup behavior when started from the command line without any parameters. Now, the connection dialog handles errors, and allows to correct the connection details and to try a new connection.

BZ#904091

This update adds the `--title STRING` option to **remote-viewer**, which makes it possible to override the default window title with user-defined text.

BZ#904094

This update adds the `--hotkeys` option that enables hotkey configuration from the command line.

BZ#905684

With this update, **virt-viewer** has been modified to show a dialog window when closing a single monitor session. This dialog asks for users' confirmation and contains an option to not be shown again.

BZ#908805

This update allows to setup a remote-viewer session from a file, for example from the RHEVM portal, via a simple browser link, without the need for a browser-specific plug-in or multi-process communication.

Users of virt-viewer are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.230. VIRT-WHO

8.230.1. RHEA-2013:1715 — virt-who enhancement update

Updated virt-who packages that add various enhancements are now available for Red Hat Enterprise Linux 6.

The virt-who packages provide an agent that collects information about virtual guests present in the system and reports them to the Red Hat Subscription Manager tool.

Enhancements

BZ#923757

With this update, support for the VMware ESXi hosts has been added to the virt-who package

BZ#1001416

This update adds support for sending information from the virt-who agent to Red Hat Satellite 5.

BZ#1002058

With this update, the virt-who package has been modified to support listing of VMware ESX virtual guest with more than 100 results.

Users of virt-who are advised to upgrade to these updated packages, which add these enhancements.

8.231. VIRTIO-WIN

8.231.1. RHBA-2013:1729 — virtio-win bug fix and enhancement update

Updated virtio-win package that fixes several bugs and adds various enhancements are now available for Red Hat Enterprise Linux 6.

The virtio-win package provides paravirtualized network drivers for most Microsoft Windows operating systems. Paravirtualized drivers are virtualization-aware drivers used by fully virtualized guests running on Red Hat Enterprise Linux. Fully virtualized guests using the paravirtualized drivers gain significantly better I/O performance than fully virtualized guests running without the drivers.

Bug Fixes

BZ#759019

Previously, when a virtio console port was hot unplugged and subsequently a new port was plugged under the same number reusing previously disowned volume groups, QEMU reported:

```
Guest moved used index from 0 to 256
```

With this update, it is no longer possible for a newly plugged port to be registered under the same number as a previously detached port. As a result, the aforementioned message is no longer displayed.

BZ#806223

Previously, when the virtio-serial device was in use, an attempt to resume the system from Sleep (S3) mode led to a stop error (the blue screen of death). The interrupt disable logic in virtio-serial has been fixed and the system can now be resumed without the stop error.

BZ#823818

Under certain circumstances, when the virtqueue was full, a guest volume group became unresponsive. Consequently, a bug check was triggered and the netkvm driver wrote to the read-only Interrupt Service Routine (ISR) register, which caused QEMU to respond with the following message:

```
virtio_ioport_write: unexpected address 0x13 value 0x0
```

This bug has been fixed and the aforementioned message is no longer displayed.

BZ#840932

Due to a missing StopIO handler in the voiserial code, a stop error (the blue screen of death) occurred while shutting down the system. With this update, the StopIO handler has been added to voiserial, and the stop error no longer occurs during the system shutdown.

BZ#856490

Due to a bug in the virtio-serial power state management code, it was impossible to turn the guest system to sleep (S3) or hibernate (S4) mode while the virtio-serial utility was in use. This bug has been fixed, and guests can now sleep or hibernate as expected.

BZ#869476

Due to a bug in the Seabios application code, the system could not be resumed from hibernation and a stop error (the blue screen of death) occurred. With this update, Seabios has been fixed and system can now be resumed as expected.

BZ#882795

Switching the system into hibernate mode (S4) while transferring data over the virtio-serial device caused a stop error (the blue screen of death). With this update, the lock acquire and release logic in virtio-serial has been changed, thus fixing this bug.

BZ#889410

In the previous release, the TCP offload code was optimized but the test-only mode of the offload emulation in drivers was not updated and therefore was not functional. Consequently, when the SW offload test mode was triggered, the network became unavailable. With this update, the SW emulation of TCP/IP offload has been removed and the guest now always applies the host offload capabilities, thus preventing the network shutdown.

BZ#902150

After resuming the system from the sleep (S3) or hibernate (S4) state, the vioserial device did not operate properly. With this update, the device handling logic in vioserial has been changed and vioserial now works properly after resuming from the S3 or S4 state.

BZ#907160

When multi-queue feature is enabled, the base address register (BAR) size of the virtio-net networks driver is increased. Previously, virtio-net was testing for BAR size during initialization and consequently failed to start. With this update, virtio-net has been modified not to test the BAR size and to use virtio features to check if additional registers can be accessed. As a result, the driver can now start without complications.

BZ#908198

Previously, the WDF installer in the virtio-win package was signed with a Red Hat signature, which overwrote the original Windows signature. With this update, WdfCoInstaller010xx.dll files are distributed without Red Hat signature.

BZ#908725

Previously, the ".inf" files in the virtio-win package were distributed without an explicit specification of revision ID, which violates certain virtio spec requirements. With this update, these ".inf" files explicitly specify revision id as part of the HW identification string.

BZ#912926, BZ#957435

Prior to this update, an interrupt could have occurred on the guest volume group during the NetKVM device shutdown. Consequently, the guest became unresponsive. With this update, the interrupts have been disabled and the device shutdown is now synchronized with DIRQL (Device Interrupt Request Level) of the device. As a result, the guest no longer hangs in the aforementioned scenario.

BZ#921200

Previously, the NetKVM driver was using the VIRTIO_ISR register for debug output in case of a guest stop error (the blue screen of death) or hangs in outgoing transfer. However, these issues should be detected by other methods. With this update, VIRTIO_ISR is no longer used for debug purposes.

BZ#950623

Previously, several debug parameters were exposed to the user. Potential change of these parameters could have influenced the functionality of the network device. This update removes these test parameters from the device manager. As a result, the debug and experimental parameters are no longer visible for the user.

BZ#950633

Due to an incorrect implementation of a corner case in a checksum testing, certain packets were dropped by the NetKVM driver. The checksum calculation has been fixed and packets are no longer dropped in the described case.

BZ#951070

Previously, certain parts of the debug information were printed without the line break, which decreased the readability of debug information. With this update, end-of-line characters have been added to the debug prints.

BZ#951481

Previously, after reinstalling the virtio-serial driver, the driver did not transfer data through the serial port. This bug has been fixed, and virtio-serial now works correctly after the reinstall sequence.

BZ#953812

Due to a bug in the write request cancellation logic of the virtio-serial driver, the guest could not be shut down while transferring data from the guest to the host over the virtio-serial port. With this update, the guest can be shut down successfully even if the data transfer from guest to host is still in progress.

BZ#955844

In virtio-win-1.6.3, Windows Server 2012 drivers were certified as the "Network-Other Device" product type, instead of "Network-LAN (Server)". Consequently, the support of "Network-LAN (Server)" was not clearly stated. This update adds the newer pre-WHQL driver build as the "Network-LAN (Server)" product type.

BZ#956290

Due to a bug in the locking sequence of the NetKVM driver, a hot unplug of the device could have caused the driver to become unresponsive. The locking sequence has been fixed and NetKVM no longer hangs after unexpected device removal.

BZ#956882

In case the drivers are set with multiple IP addressees, an incorrect parsing of the IP address structure in the driver could have caused a stop error (the blue screen of death). As this parsing was done for debug purposes only, with this update it has been completely removed from the drivers.

BZ#957505, BZ#960503

Due to a race condition, setting the virtio features to "0" before VM hibernation could trigger a QEMU assertion or a network interruption. With this update, these features are no longer set to "0", thus preventing both assertion and network failure in the aforementioned scenario.

BZ#966809

Previously, resuming the system from hibernation mode (S4) after hot plugging a virtio-serial device led to a stop error (the blue screen of death). This bug has been fixed and guest can resume successfully in the described case.

BZ#972310

In case of failed initialization of the NetKVM driver, certain internal system data structures were also not initialized. Consequently, access to uninitialized data structures during driver unload caused a stop error (the blue screen of death). With this update, the partial driver initialization is handled correctly and a stop error no longer occurs.

BZ#982940

Previously, the INF file incorrectly indicated 1GB connection rate. While the connection rate of this size could be set from the device manager, it was impossible to set 1GB rate using a command-line tool. This update sets the correct connection rate in INF file and this rate can now be configured with the command line as expected.

BZ#988302

An automatic conversion to Visual Studio projects set an incorrect calling convention for 32bit binaries of the netsh plug-in. Consequently, the command-line options of the netsh plug-in did not work on 32bit operation systems. With this update, the calling convention has been changed and the netsh plug-in command-line options now work correctly on 32bit architectures.

Enhancements**BZ#713130**

With this update, support for the event index feature has been added. This feature reduces CPU utilization per megabyte for most operations, for example during the network transfer.

BZ#904934

A new mechanism of setting the MAC address was introduced in QEMU. Setting the MAC address can now be made as an atomic change, and QEMU is able to test MAC address for validity.

BZ#920011

By default, the OVS switch removes the priority header from packets, so the priority and VLAN tags are stripped from the packets originated in Windows guests. This behavior is now configurable with

the "other-config:priority-tags" option. To preserve the priority tags, for each port added to the OVS bridge the following command needs to be executed:

```
ovs-vsctl set port <PORT_NAME> other-config:priority-tags=true
```

BZ#948017

This update adds support for VSS (Volume Shadow Copy Service) to the qemu-ga-win agent. VSS is a Microsoft Windows API that allows for consistent freeze and thaw operations. With this feature, snapshots taken while the virtual machine is running are consistent through the whole stack from the block layer to the guest applications and can be used for backup purposes.

BZ#950424

With this update, the virtio indirect buffers feature can be used in the guest driver, providing enhanced virtio capability to transfer fragmented network traffic.

BZ#950509

With this update, NetKVM driver build facilities have been converted to Visual Studio projects or solutions, providing compatibility with the with Windows 8 DDK.

BZ#950617

This update implements RSS (Receive Side Scaling), which improves RX performance on Windows SMP guests.

BZ#971141

The Windows guest agent is now fully supported and delivered with its own installer in the Supplementary channel together with the virtio-win drivers.

Users of virtio-win are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements.

8.232. WATCHDOG

8.232.1. RHBA-2013:1656 — watchdog bug fix update

Updated watchdog package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The watchdog package provides a user-space application which can be configured to provide updates to a hardware or software watchdog timer via the Linux kernel's watchdog interface.



NOTE

The watchdog package has been upgraded to upstream version 5.6, which provides one bug fix over the previous version. The previous version failed to reset the lpmi Linux watchdog timer on a Sun Fire X4100 server. Consequently, the system started to reboot continuously if the watchdog daemon was started. A patch has been applied to address this bug, and resetting the watchdog timer now works as expected on Sun Fire X4100. (BZ#870217)

Users of watchdog are advised to upgrade to this updated package, which fixes this bug.

8.233. WEBKITGTK

8.233.1. RHBA-2013:1613 — webkitgtk bug fix update

Updated webkitgtk packages that fix one bug are now available for Red Hat Enterprise Linux 6.

WebKitGTK+ is the port of the portable web rendering engine WebKit to the GTK+ platform.

Bug Fix

BZ#966571

Previously, the just-in-time (JIT) compilation was not enabled for JavaScript in WebKitGTK+. As a consequence, some websites experienced problems with loading and performance. This update enables the JIT compilation for 32-bit Intel, Intel 64, and AMD64 architectures and the unwanted behavior no longer occurs.

Users of webkitgtk are advised to upgrade to these updated packages, which fix this bug.

8.234. WIRESHARK

8.234.1. RHSA-2013:1569 — Moderate: wireshark security, bug fix, and enhancement update

Updated wireshark packages that fix multiple security issues, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

Wireshark, previously known as Ethereal, is a network protocol analyzer. It is used to capture and browse the traffic running on a computer network.

Security Fixes

CVE-2013-3559, CVE-2013-4083

Two flaws were found in Wireshark. If Wireshark read a malformed packet off a network or opened a malicious dump file, it could crash or, possibly, execute arbitrary code as the user running Wireshark.

[CVE-2012-2392](#), [CVE-2012-3825](#), [CVE-2012-4285](#), [CVE-2012-4288](#), [CVE-2012-4289](#), [CVE-2012-4290](#), [CVE-2012-4291](#), [CVE-2012-4292](#), [CVE-2012-5595](#), [CVE-2012-5597](#), [CVE-2012-5598](#), [CVE-2012-5599](#), [CVE-2012-5600](#), [CVE-2012-6056](#), [CVE-2012-6059](#), [CVE-2012-6060](#), [CVE-2012-6061](#), [CVE-2012-6062](#), [CVE-2013-3557](#), [CVE-2013-3561](#), [CVE-2013-4081](#), [CVE-2013-4927](#), [CVE-2013-4931](#), [CVE-2013-4932](#), [CVE-2013-4933](#), [CVE-2013-4934](#), [CVE-2013-4935](#), [CVE-2013-4936](#), [CVE-2013-5721](#)

Several denial of service flaws were found in Wireshark. Wireshark could crash or stop responding if it read a malformed packet off a network, or opened a malicious dump file.



NOTE

The Wireshark packages have been upgraded to upstream version 1.8.10, which provides a number of bug fixes and enhancements over the previous versions. For more information on the bugs fixed, enhancements included, and supported protocols introduced, refer to the Wireshark Release Notes:

<http://www.wireshark.org/docs/relnotes/wireshark-1.8.0.html>

<http://www.wireshark.org/docs/relnotes/wireshark-1.6.0.html>

<http://www.wireshark.org/docs/relnotes/wireshark-1.4.0.html>

(BZ#711024)

Bug Fixes

BZ#750712

Previously, Wireshark did not parse the RECLAIM-COMPLETE opcode when inspecting traffic generated by NFSv4.1. A patch has been provided to enable the parsing of the RECLAIM_COMPLETE opcode, and Wireshark is now able to properly dissect and handle NFSv4.1 traffic.

BZ#832021

Prior to this update, frame arrival times in a text file were reported one hour ahead from the timestamps in the packet capture file. This resulted in various failures being reported by the `dfilter-test.py` test suite. To fix this bug, frame arrival timestamps have been shifted by one hour, thus fixing this bug.

BZ#1004636

The `"tshark -D"` command returned output to `STDERR` instead of `STDOUT`, which could break scripts that are parsing the `"tshark -D"` output. This bug has been fixed, and the `"tshark -D"` command now writes output data to a correct standard stream.

BZ#715560

Due to an array overrun, Wireshark could experience undefined program behavior or could unexpectedly terminate. With this update, proper array handling ensures Wireshark no longer crashes in the described scenario.

BZ#659661

Previously, the `dfstest` and `randpkt` command line utilities lacked manual pages. This update adds proper manual pages for both utilities.

Enhancement

BZ#699636, BZ#858976

With this update, Wireshark is able to properly dissect and handle InfiniBand and GlusterFS traffic.

All Wireshark users are advised to upgrade to these updated packages, which contain backported patches to correct these issues and add these enhancements. All running instances of Wireshark must be restarted for the update to take effect.

8.235. XFSPROGS

8.235.1. RHBA-2013:1657 — xfsprogs bug fix update

Updated xfsprogs packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The xfsprogs packages contain a set of commands to use the XFS file system, including the `mkfs.xfs` command to construct an XFS file system.

Bug Fixes

BZ#893904

Previously, when the `ag_stride` option was specified as an option to the `xfs_repair` command, the command terminated with a segmentation fault. This update modifies the underlying XFS repair code so the `ag_stride` option is now properly handled.

BZ#950691

Previously on certain file systems, the `xfs_repair` utility could occasionally emit warnings such as "7f61c041b700: Badness in key lookup (length)". These were harmless errors related to cache lookup failures, however these messages might have concerned users running `xfs_repair`. This update ensures that these warnings are no longer encountered.

BZ#961501

When stripe geometry was specified manually to the `mkfs.xfs` command, `mkfs.xfs` did not properly select "multidisk mode" as it does when stripe geometry is automatically detected. As a result, a less than optimal number of allocation groups was created. With this update, multidisk mode is selected properly, and a larger number of allocation groups is created.

BZ#962394

The `xfs_io(8)` man page did not contain documentation for the `chproj`, `lsproj`, and `setfl` commands. These commands are now documented in the man page. Also, the man page has been corrected to refer to section 2 of the `fallocate()` system call.

BZ#962397

Previously, the `xfs_logprint` command could abort with the "xlog_print_trans_inode: illegal inode type" error when it encountered a multiply-logged inode field. This update modifies the underlying logprint code so that multiply-logged inode fields are now handled properly, and `xfs_logprint` completes successfully as expected.

BZ#964216

Previously, the `xfs_repair` utility was unable to properly handle fragmented multi-block version 2 directories, which could, under certain circumstances, result in an `xfs_repair` segmentation fault. This update modifies the underlying code so `xfs_repair` can now operate on fragmented version 2 directories as expected.

BZ#987538

The `mkfs.xfs(8)` man page did not contain a description of the "`-d noalign`" option, which disables automatic storage geometry detection at `mkfs` time. With this update, the option is now properly documented.

BZ#1002908

Previously, the `xfs_logprint` command could fail when encountering a continued inode transaction or a wrapped log. This update modifies the underlying logprint code so that continued inode transactions and wrapped logs are now handled properly, and `xfs_logprint` completes successfully as expected.

Users of `xfprogs` are advised to upgrade to these updated packages, which fix these bugs.

8.236. XMLRPC-C

8.236.1. RHBA-2013:1254 — `xmlrpc-c` bug fix update

Updated `xmlrpc-c` packages that fix one bug are now available.

XML-RPC is a remote procedure call (RPC) protocol that uses XML to encode its calls and HTTP as a transport mechanism.

Bug Fix

BZ#809819

Previously, features listed when the `--help` command was run were not consistent with the list when the `--features` command was run. Also, running the reproducer script resulted in "Unrecognized token" errors. With this update, listed features are consistent, and "Unrecognized token" errors are no longer displayed.

Users of `xmlrpc-c` are advised to upgrade to these updated packages, which fix this bug.

8.237. XORG-X11-DRV-ATI

8.237.1. RHBA-2013:1597 — `xorg-x11-drv-ati` bug fix and enhancement update

Updated `xorg-x11-drv-ati` packages that fix one bug and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-ati` packages provide a driver for ATI graphics cards for the X.Org implementation of the X Window System.

Bug Fix

BZ#907616

Previously, the hot plug detection and mode setting were not performed automatically. As a consequence, the user had to run the utility program to enforce the changes when plugging in multiple monitors. The driver has been updated, and the hotplug detection and mode setting now work as expected.

Enhancements

BZ#795919

Support for new graphics processing unit (GPU) hardware has been added.

BZ#822280, BZ#879102, BZ#882086

Support for hardware acceleration, including OpenGL, of the Radeon HD 7000 Series has been added.

Users of `xorg-x11-drv-ati` are advised to upgrade to these updated packages, which fix this bug and add these enhancements.

8.238. XORG-X11-DRV-INTEL

8.238.1. RHBA-2013:1589 — `xorg-x11-drv-intel` bug fix and enhancement update

Updated `xorg-x11-drv-intel` packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-intel` package contains an Intel integrated graphics video driver for the X.Org implementation of the X Window System.



NOTE

Within the `xorg-x11` package, `intel-gpu-tools` has been upgraded to upstream version 2.21.12, which provides a number of bug fixes and enhancements over the previous version. (BZ#[906036](#))

Bug Fixes

BZ#[886191](#)

Previously, when Red Hat Enterprise Linux 6 was installed on some workstations, Anaconda chose the `fbdev` driver instead of the `i915` Intel driver. With this update, the Intel driver is selected, thus fixing the bug.

BZ#[999334](#)

When building rpms from src rpms, the spec definition and sources produced files, which were not packaged in the last phase of the `rpm` build procedure. Consequently, `rpm` build did not create rpm(s) even if everything else was built successfully, and the following error message was returned:

```
Installed (but unpackaged) file(s) found
```

This bug has been fixed and all files are now packaged as expected.

Users of `xorg-x11-drv-intel` are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

8.239. XORG-X11-DRV-MGA

8.239.1. RHBA-2013:1610 — `xorg-x11-drv-mga` bug fix update

Updated `xorg-x11-drv-mga` packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-mga` packages provide a video driver for Matrox G-series chipsets for the X.Org implementation of the X Window System.

Bug Fixes

BZ#894959

Prior to this update, the graphical user interface could appear distorted on 19-inch monitors with the 16:9 ratio. The `xorg-x11-drv-mga` packages have been fixed, and so the distortion no longer occurs in this scenario.

BZ#918017

Previously, resolutions higher than 1440x900 were not available with Red Hat Enterprise Linux 6.4 using the MGA G200e chips. Consequently, the Matrox driver did not allow native resolutions to be reached for many monitors. With this update, the X Server no longer discards larger resolution modes, and resolutions higher than 1440x900 are now available.

Users of `xorg-x11-drv-mga` are advised to upgrade to these updated packages, which fix these bugs.

8.240. XORG-X11-DRV-NOUVEAU

8.240.1. RHBA-2013:1664 — [xorg-x11-drv-nouveau bug fix update](#)

Updated `xorg-x11-drv-nouveau` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-nouveau` packages provide the X.Org X11 nouveau video driver for NVIDIA graphics chipsets.

Bug Fix

BZ#876566

Previously, when using a VGA-compatible controller for certain NVIDIA Quadro graphics cards, the rendercheck test suite was not able to perform the complete check due to rendering problems. The `xorg-x11-drv-nouveau` packages have been fixed, rendering problems no longer occur, and the test suite completes the check as expected.

Users of `xorg-x11-drv-nouveau` are advised to upgrade to these updated packages, which fix this bug.

8.241. XORG-X11-DRV-QXL

8.241.1. RHBA-2013:1650 — [xorg-x11-drv-qxl bug fix update](#)

Updated `xorg-x11-drv-qxl` packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-qxl` packages provide an X11 video driver for the QEMU QXL video accelerator. This driver makes it possible to use Red Hat Enterprise Linux 6 as a guest operating system under the KVM kernel module and the QEMU multi-platform emulator, using the SPICE protocol.

Bug Fixes

BZ#929037

When the user tried to start a guest with Red Hat Enterprise Linux 6 on a host with Red Hat Enterprise Linux 5, the QEMU QXL video accelerator failed with a segmentation fault. As a

consequence, the guest was not able to start the system GUI. This update applies a patch to fix this bug and the guest now starts correctly.

BZ#951000

When using multiple QXL devices with the Xinerama extension, or multiple QXL devices while each being a separate screen, an attempt to set a resolution higher than 1024 x 768 pixels in the `xorg.conf` file failed with an error. With this update, the underlying source code has been modified and the resolution can now be set as expected.

Users of `xorg-x11-drv-qxl` are advised to upgrade to these updated packages, which fix these bugs.

8.242. XORG-X11-DRV-SYNAPTICS

8.242.1. RHBA-2013:1644 — [xorg-x11-drv-synaptics bug fix update](#)

Updated `xorg-x11-drv-synaptics` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-synaptics` packages contain the X.Org X11 input drivers for Synaptics touchpads.

Bug Fix

BZ#988174

Previously, the synaptics driver scaled input coordinates based on the device resolution. Consequently, the X server could not apply the uniform resolution-based scaling for other devices in relative mode. The synaptics driver has been fixed to apply the scaling feature only if the X server does not support the per-device resolution scaling.

Users of `xorg-x11-drv-synaptics` are advised to upgrade to these updated packages, which fix this bug.

8.243. XORG-X11-DRV-WACOM

8.243.1. RHBA-2013:1568 — [xorg-x11-drv-wacom bug fix update](#)

Updated `xorg-x11-drv-wacom` packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The `xorg-x11-drv-wacom` packages provide an X Window System input device driver that allows the X server to handle Wacom tablets with extended functionality.

Bug Fix

BZ#920385

Previously, the `xsetwacom` command was unable to map the Esc key to one of the buttons on a Wacom Cintiq 22HD tablet using the `xsetwacom` command. As a consequence, the command failed without displaying an error message, and the key was not mapped. This bug has now been fixed, and, as a result, the driver now maps the Esc key as expected.

Users of `xorg-x11-drv-wacom` are advised to upgrade to these updated packages, which fix this bug.

8.244. XORG-X11-SERVER

8.244.1. RHSA-2013:1620 — Low: xorg-x11-server security and bug fix update

Updated xorg-x11-server packages that fix one security issue and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having low security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

X.Org is an open source implementation of the X Window System. It provides the basic low-level functionality that full-fledged graphical user interfaces are designed upon.

Security Fix

CVE-2013-1940

A flaw was found in the way the X.org X11 server registered new hot plugged devices. If a local user switched to a different session and plugged in a new device, input from that device could become available in the previous session, possibly leading to information disclosure.

This issue was found by David Airlie and Peter Hutterer of Red Hat.

Bug Fixes

BZ#915202

A previous upstream patch modified the Xephyr X server to be resizeable, however, it did not enable the resize functionality by default. As a consequence, X sandboxes were not resizeable on Red Hat Enterprise Linux 6.4 and later. This update enables the resize functionality by default so that X sandboxes can now be resized as expected.

BZ#957298

In Red Hat Enterprise Linux 6, the X Security extension (XC-SECURITY) has been disabled and replaced by X Access Control Extension (XACE). However, XACE does not yet include functionality that was previously available in XC-SECURITY. With this update, XC-SECURITY is enabled in the xorg-x11-server spec file on Red Hat Enterprise Linux 6.

BZ#969538

Upstream code changes to extension initialization accidentally disabled the GLX extension in Xvfb (the X virtual frame buffer), rendering headless 3D applications not functional. An upstream patch to this problem has been backported so the GLX extension is enabled again, and applications relying on this extension work as expected.

All xorg-x11-server users are advised to upgrade to these updated packages, which contain backported patches to correct these issues.

8.245. XORG-X11-XINIT

8.245.1. RHBA-2013:1538 — xorg-x11-xinit bug fix update

Updated xorg-x11-xinit packages that fix one bug are now available for Red Hat Enterprise Linux 6.

X.Org is an open source implementation of the X Window System providing basic low-level functionality that full-fledged desktop environments such as GNOME and KDE are built on top of. The xorg-x11-xinit packages contain the X.Org X Window System xinit startup scripts.

Bug Fix

BZ#811289

Previously, the startx script did not handle the xserverrc file properly. If the xserverrc file existed in the /etc/X11/xinit/ directory, the script failed with the following error message:

```
Fatal server error: Unrecognized option: /etc/X11/xinit/xserverrc
```

With this update, the X session is started using options from the xserverrc file, and the startx script now properly handles the xserverrc file.

Users of xorg-x11-xinit are advised to upgrade to these updated packages, which fix this bug.

8.246. YABOOT

8.246.1. RHBA-2013:1561 — yaboot bug fix and enhancement update

Updated yaboot packages that fix two bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The yaboot packages provide a boot loader for Open Firmware based PowerPC systems. Yaboot can be used to boot IBM eServer System p machines.

Bug Fixes

BZ#903855

Previously, the client was overwriting the gateway IP address for the Trivial File Transfer Protocol (TFTP) file transfer. When installing through the network using VLAN tags, the boot failed when the server was in a different IP subnetwork. This update ensures that the rest of the parameter strip can be parsed correctly, and failures no longer occur in the aforementioned scenario.

BZ#968046

As there was not enough room between the first allocation and the bottom of the firmware, user attempts to load a ramdisk failed when the firmware was at 32MB (0200000). This update adds the ability to be able to determine how big the initrd memory will be yaboot . As a result, yaboot can accurately place a buffer in the memory, and ramdisk load failures no longer occur.

Enhancement

BZ#947101

This update adds GUID Partition Table (GPT) support to yaboot, because previously yaboot supported the DOS partition format which has a limit of up to 2TB for 512B sectors. So even if there are large disks, this limit forces users to format all devices to 2TB. With GPT support in yaboot, users can now use larger disks.

Users of yaboot are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

8.247. YUM-RHN-PLUGIN

8.247.1. RHBA-2013:1086 — yum-rhn-plugin bug fix update

Updated yum-rhn-plugin packages that fix one bug are now available.

The yum-rhn-plugin package provides support for connecting to Red Hat Network (RHN). Systems registered with RHN are able to update and install packages from Red Hat Network.

Bug Fix

BZ#949649

The RHN Proxy did not work properly if separated from a parent by a slow enough network. Consequently, users who attempted to download larger repodata files and RPMs experienced timeouts. This update changes both RHN Proxy and Red Hat Enterprise Linux RHN Client to allow all communications to obey a configured timeout value for connections.

Users of yum-rhn-plugin are advised to upgrade to these updated packages, which fix this bug.

8.247.2. RHBA-2013:1703 — yum-rhn-plugin bug fix update

Updated yum-rhn-plugin packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The yum-rhn-plugin packages allow the Yum package manager to access content from Red Hat Network.

Bug Fix

BZ#960524, BZ#988895

Prior to this update, an attempt to install an already-installed package led to an empty transaction that was incorrectly identified as an error. Consequently, the yum-rhn-plugin reported a failed installation action. With this update, yum-rhn-plugin has been modified to return a success code for empty transactions. As a result, a successful installation action is now reported when the package is already installed.

Users of yum-rhn-plugin are advised to upgrade to these updated packages, which fix this bug.

8.248. ZSH

8.248.1. RHEA-2013:1557 — zsh enhancement update

Updated zsh packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

The zsh shell is a command interpreter which can be used as an interactive login shell and as a shell script command processor. Zsh resembles the ksh shell (the Korn shell), but includes many enhancements. Zsh supports command line editing, built-in spelling correction, programmable command completion, shell functions (with autoloading), a history mechanism, and more.

Enhancement

BZ#820530

This update adds the `${NAME:OFFSET:LENGTH}` substitution to zsh, which enables users to use the syntax `${NAME:OFFSET:LENGTH}` to display a substring.

Users of zsh are advised to upgrade to these updated packages, which add this enhancement.

APPENDIX A. REVISION HISTORY

Revision 1-0.48	Thu Jan 21 2016	Lenka Špačková
Added information about the removed <code>systemtap-grapher</code> package to the Deprecated Functionality Chapter.		
Revision 1-0.47	Tue Jan 27 2015	Milan Navrátil
Updated the Red Hat Enterprise Linux 6.5 Technical Notes.		
Revision 1-0.42	Mon Oct 20 2014	Miroslav Svoboda
Updated the Red Hat Enterprise Linux 6.5 Technical Notes with the latest kernel changes.		
Revision 1-0.32	Fri Jun 20 2014	Miroslav Svoboda
Updated the Red Hat Enterprise Linux 6.5 Technical Notes with the latest kernel erratum, RHSA-2014-0771.		
Revision 1-0.30	Mon Jun 02 2014	Eliška Slobodová
Clarified that iSCSI and FCoE boot are fully supported features in Red Hat Enterprise Linux 6.5.		
Revision 1-0.29	Thu May 08 2014	Miroslav Svoboda
Updated the Red Hat Enterprise Linux 6.5 Technical Notes with the latest kernel erratum, RHSA-2014-0475.		
Revision 1-0.25	Tue Feb 18 2014	Eliška Slobodová
Added a Mellanox SR-IOV Technology Preview.		
Revision 1-0.23	Wed Feb 12 2014	Miroslav Svoboda
Updated the Red Hat Enterprise Linux 6.5 Technical Notes with the latest kernel erratum, RHSA-2014-0159.		
Revision 1-0.22	Wed Jan 22 2014	Eliška Slobodová
Added the missing eCryptfs Technology Preview.		
Revision 1-0.21	Fri Jan 10 2014	Eliška Slobodová
Fixed several typos.		
Revision 1-0.17	Fri Dec 13 2013	Miroslav Svoboda
Updated the Red Hat Enterprise Linux 6.5 Technical Notes with the latest kernel erratum, RHSA-2013-1801.		
Revision 1-0.15	Thu Nov 21 2013	Eliška Slobodová
Release of the Red Hat Enterprise Linux 6.5 Technical Notes.		
Revision 1-0.0	Thu Oct 03 2013	Eliška Slobodová
Release of the Red Hat Enterprise Linux 6.5 Beta Technical Notes.		