

Red Hat Enterprise Linux 6

6.3 Technical Notes

Detailed notes on the changes implemented in Red Hat Enterprise Linux 6.3

Edition 3

Last Updated: 2025-01-15

Red Hat Enterprise Linux 6 6.3 Technical Notes

Detailed notes on the changes implemented in Red Hat Enterprise Linux 6.3
Edition 3

Red Hat Engineering Content Services

Legal Notice

Copyright © 2012 Red Hat, Inc.

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

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

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat 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 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 6.3 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 6.2 and minor release Red Hat Enterprise Linux 6.3.

Table of Contents

PREFACE 10

CHAPTER 1. IMPORTANT CHANGES TO EXTERNAL KERNEL PARAMETERS 11

CHAPTER 2. TECHNOLOGY PREVIEWS 14
2.1. STORAGE AND FILE SYSTEMS 14
2.2. NETWORKING 17
2.3. CLUSTERING AND HIGH AVAILABILITY 17
2.4. AUTHENTICATION 18
2.5. SECURITY 19
2.6. DEVICES 19
2.7. KERNEL 20
2.8. VIRTUALIZATION 21
2.9. RESOURCE MANAGEMENT 23

CHAPTER 3. KNOWN ISSUES 24
3.1. INSTALLATION 24
3.2. ENTITLEMENT 26
3.3. DEPLOYMENT 26
3.4. VIRTUALIZATION 27
3.5. STORAGE AND FILE SYSTEMS 29
3.6. NETWORKING 30
3.7. CLUSTERING 31
3.8. AUTHENTICATION 31
3.9. DEVICES 35
3.10. KERNEL 37
3.11. DESKTOP 43
3.12. TOOLS 44

CHAPTER 4. NEW PACKAGES 47
4.1. RHEA-2012:0842 – NEW PACKAGE: BYZANZ 47
4.2. RHEA-2012:0797 – NEW PACKAGES: CRASH-GCORE-COMMAND 47
4.3. RHEA-2012:0831 – NEW PACKAGE: DEVICE-MAPPER-PERSISTENT-DATA 47
4.4. RHEA-2012:0814 – NEW PACKAGE: I2C-TOOLS 47
4.5. RHEA-2012:0829 – NEW PACKAGES: IPSET AND LIBMNL 48
4.6. RHEA-2012:0840 – NEW PACKAGES: JAVA-1.7.0-IBM 48
4.7. RHEA-2012:0981 – NEW PACKAGES: JAVA-1.7.0-OPENJDK 48
4.8. RHEA-2012:0838 – NEW PACKAGE: JAVA-1.7.0-ORACLE 49
4.9. RHEA-2012:1038 – NEW PACKAGES: KMOD-BNX2X, KMOD-BNX2, KMOD-BNX2I, KMOD-BNX2FC 49
4.10. RHEA-2012:1576 – NEW PACKAGES: KMOD-PCH_GBE 50
4.11. RHEA-2012:0825 – NEW PACKAGE: LEDMON 50
4.12. RHEA-2012:0812 – NEW PACKAGE: LIBQB 50
4.13. RHEA-2012:0798 – NEW PACKAGES: LIBREOFFICE 51
4.14. RHEA-2012:0868 – NEW PACKAGES: LIBWACOM 51
4.15. RHEA-2012:0890 – NEW PACKAGE: NUMAD 51
4.16. RHEA-2012:0826 – NEW PACKAGE: PPC64-DIAG 51
4.17. RHEA-2012:0806 – NEW PACKAGES: SCL-UTILS 52
4.18. RHEA-2012:0823 – NEW PACKAGE: SUBSCRIPTION-MANAGER-MIGRATION-DATA 52
4.19. RHEA-2012:0853 – NEW PACKAGES: USBREDIR 52
4.20. RHEA-2012:0965 – NEW PACKAGE: VIRT-P2V 52
4.21. RHEA-2012:0786 – NEW PACKAGES: KMOD-HPSA 53

CHAPTER 5. PACKAGE UPDATES 54
5.1. 389-DS-BASE 54
5.2. ABRT AND LIBREPORT 63
5.3. ABRT, LIBREPORT, BTPARSER, AND PYTHON-MEH 64
5.4. ACROREAD 66
5.5. ALSA-UTILS 67
5.6. ANACONDA 67
5.7. ATLAS 74
5.8. AUDIT 75
5.9. AUGEAS 76
5.10. AUTHCONFIG 77
5.11. AUTOFS 78
5.12. AXIS 79
5.13. BACULA 80
5.14. BIND-DYNDB-LDAP 81
5.15. BIND 83
5.16. BINUTILS 88
5.17. BIOSDEVNAME 89
5.18. BRLTTY 89
5.19. BUSYBOX 90
5.20. BYACC 91
5.21. C-ARES 91
5.22. CDRKIT 92
5.23. CERTMONGER 92
5.24. CHKCONFIG 93
5.25. CIFS-UTILS 94
5.26. CLUSTER AND GFS2-UTILS 96
5.27. CLUSTER-GLUE 98
5.28. CLUSTERMON 98
5.29. CLUSTER 99
5.30. CONMAN 100
5.31. CONTROL-CENTER 101
5.32. COOLKEY 101
5.33. COREUTILS 102
5.34. COROSYNC 103
5.35. CPIO 105
5.36. CPUSPEED 106
5.37. CRASH 107
5.38. CRASH-TRACE-COMMAND 108
5.39. CREATEREPO 109
5.40. CRYPTSETUP-LUKS 109
5.41. CTDB 109
5.42. CUPS 110
5.43. CVS 111
5.44. CYRUS-SASL 112
5.45. DASH 112
5.46. DB4 113
5.47. DBUS 114
5.48. DEVICE-MAPPER-MULTIPATH 114
5.49. DHCP 117
5.50. DING-LIBS 119
5.51. DMRAID 120
5.52. DNSMASQ 120

5.53. DOCBOOK-UTILS 121
5.54. DRACUT 121
5.55. DROPWATCH 123
5.56. DVD+RW-TOOLS 124
5.57. E2FSPROGS 124
5.58. EFIBOOTMGR 125
5.59. ELINKS 125
5.60. ESPEAK 126
5.61. EXPECT 126
5.62. FCOE-TARGET-UTILS 127
5.63. FCOE-UTILS 127
5.64. FEBOOTSTRAP 128
5.65. FENCE-AGENTS 128
5.66. FENCE-VIRT 131
5.67. FILE 131
5.68. FIREFOX 133
5.69. FIRSTBOOT 141
5.70. FLASH-PLUGIN 141
5.71. FONTFORGE 145
5.72. FPRINTD 145
5.73. FREERADIUS 146
5.74. FREETYPE 147
5.75. FTP 148
5.76. GAWK 150
5.77. GCC 151
5.78. GDB 152
5.79. GDM 153
5.80. GD 153
5.81. GEGL 154
5.82. GERONIMO-SPECS 154
5.83. GHOSTSCRIPT 154
5.84. GIMP 155
5.85. GLIB2 156
5.86. GLIBC 157
5.87. GNOME-DESKTOP 162
5.88. GNOME-KEYRING 163
5.89. GNOME-PACKAGEKIT 163
5.90. GNOME-POWER-MANAGER 164
5.91. GNOME-SCREENSAVER 164
5.92. GNOME-SETTINGS-DAEMON 165
5.93. GNOME-SYSTEM-MONITOR 166
5.94. GNOME-TERMINAL 166
5.95. GRAPHVIZ 167
5.96. GREP 167
5.97. GRUBBY 168
5.98. GRUB 168
5.99. GSTREAMER-PLUGINS-BASE 169
5.100. GTK2 169
5.101. GVFS 170
5.102. HIVEX 171
5.103. HSQLDB 171
5.104. HWDATA 171
5.105. ICEDTEA-WEB 172

5.106. IMSETTINGS 173
5.107. INDENT 173
5.108. INITSCRIPTS 174
5.109. IOK 177
5.110. IPA 178
5.111. IPMITOOL 185
5.112. IPROUTE 186
5.113. IPRUTILS 187
5.114. IPTRAF 188
5.115. IPVSADM 188
5.116. IRQBALANCE 189
5.117. IRSSI 189
5.118. ISCSI-INITIATOR-UTILS 190
5.119. JAKARTA-COMMONS-HTTPCLIENT 191
5.120. JAVA-1.5.0-IBM 191
5.121. JAVA-1.6.0-IBM 192
5.122. JAVA-1.6.0-OPENJDK 193
5.123. JAVA-1.6.0-SUN 198
5.124. JAVA-1.7.0-IBM 199
5.125. JAVA-1.7.0-OPENJDK 200
5.126. JAVA-1.7.0-ORACLE 205
5.127. JSS 207
5.128. KABI-WHITELISTS 208
5.129. KDEARTWORK 209
5.130. KDEBASE 209
5.131. KDEBASE-WORKSPACE 210
5.132. KDELIBS3 210
5.133. KDELIBS 211
5.134. KDEPIM 213
5.135. KERNEL 213
5.136. KEXEC-TOOLS 261
5.137. KEYUTILS 266
5.138. KRB5 266
5.139. KSH 269
5.140. LATENCYTOP 270
5.141. LIBBONOBO 271
5.142. LIBBURN 271
5.143. LIBCGROUP 272
5.144. LIBDVDREAD 272
5.145. LIBERATION-FONTS 272
5.146. LIBEVENT 273
5.147. LIBEXIF 273
5.148. LIBGUESTFS 274
5.149. LIBGWEATHER 276
5.150. LIBHBAAPI 276
5.151. LIBHBALINUX 277
5.152. LIBIBVERBS-ROCEE AND LIBMLX4-ROCEE 277
5.153. LIBPROXY 278
5.154. LIBREOFFICE 278
5.155. LIBSELINUX 279
5.156. LIBSERVICELOG 279
5.157. LIBSSH2 280
5.158. LIBTAR 280

5.159. LIBTIFF 281
5.160. LIBUNISTRING 282
5.161. LIBUSB1 282
5.162. LIBUSER 283
5.163. LIBVIRT-CIM 283
5.164. LIBVIRT-JAVA 284
5.165. LIBVIRT-QMF 284
5.166. LIBVIRT 285
5.167. LIBXKLAVIER 299
5.168. LIBXML2 299
5.169. LIBXSLT 300
5.170. LLDPAD 301
5.171. LM_SENSORS 303
5.172. LOGROTATE 303
5.173. LOHIT-KANNADA-FONTS 304
5.174. LOHIT-TELUGU-FONTS 305
5.175. LSOF 305
5.176. LSVPD 305
5.177. LTRACE 306
5.178. LUCI 307
5.179. LVM2 309
5.180. M2CRYPTO 316
5.181. MAILMAN 316
5.182. MAKE 317
5.183. MAN-PAGES-FR 317
5.184. MAN-PAGES-OVERRIDES 318
5.185. MAN 319
5.186. MATAHARI 320
5.187. MCELOG 321
5.188. MDADM 322
5.189. METACITY 324
5.190. MICROCODE_CTL 325
5.191. MINGW32-LIBXML2 325
5.192. MINGW32-MATAHARI 327
5.193. MINGW32-QPID-CPP 327
5.194. MKBOOTDISK 328
5.195. MLOCATE 328
5.196. MOD_AUTH_KERB 329
5.197. MOD_AUTHZ_LDAP 329
5.198. MOD_NSS 330
5.199. MODULE-INIT-TOOLS 330
5.200. MOD_WSGI 331
5.201. MRTG 331
5.202. MT-ST 332
5.203. MYSQL-CONNECTOR-JAVA 332
5.204. MYSQL 333
5.205. NAUTILUS 335
5.206. NET-SNMP 336
5.207. NETWORKMANAGER-OPENSWAN 340
5.208. NETWORKMANAGER 340
5.209. NFS4-ACL-TOOLS 343
5.210. NFS-UTILS 343
5.211. NMAP 345

5.212. NSPLUGINWRAPPER 345
5.213. NSS, NSS-UTIL, AND NSPR 346
5.214. NSS-PAM-LDAPD 349
5.215. NSS 349
5.216. NUMACTL 350
5.217. NUMPY 350
5.218. OPENJPEG 351
5.219. OPENLDAP 352
5.220. OPENMOTIF 355
5.221. OPENSSH 355
5.222. OPENSSL 357
5.223. OPENSWAN 358
5.224. OPROFILE 362
5.225. ORBIT2 362
5.226. PACEMAKER 363
5.227. PACKAGEKIT 364
5.228. PAM_PKCS11 365
5.229. PANGO 365
5.230. PARTED 365
5.231. PCRE 366
5.232. PCSC-LITE 367
5.233. PERL-DBD-PG 368
5.234. PERL-GSSAPI 368
5.235. PERL-IPC-RUN3 369
5.236. PERL-IPC-RUN 369
5.237. PERL-SOAP-LITE 369
5.238. PERL-SYS-VIRT 370
5.239. PERL 371
5.240. PHP-PECL-APC 372
5.241. PHP-PECL-MEMCACHE 373
5.242. PHP 373
5.243. PIDGIN 375
5.244. PIRANHA 376
5.245. PKI-CORE 377
5.246. PM-UTILS 378
5.247. POLICYCOREUTILS 379
5.248. PORTRESERVE 380
5.249. POSTGRESQL AND POSTGRESQL84 380
5.250. POSTGRESQL-JDBC 382
5.251. PPC64-UTILS 382
5.252. PROCPS 382
5.253. PSACCT 383
5.254. PULSEAUDIO 384
5.255. PYKICKSTART 385
5.256. PYQT4 385
5.257. PYTHON-CONFIGSHELL 386
5.258. PYTHON-MEMCACHED 386
5.259. PYTHON-PASTE-SCRIPT 387
5.260. PYTHON-REPOZE-WHO 387
5.261. PYTHON-RHSM 387
5.262. PYTHON-RTSLIB 388
5.263. PYTHON 389
5.264. PYTHON-VIRTINST 389

5.265. QEMU-KVM 391
5.266. QL2400-FIRMWARE 398
5.267. QL2500-FIRMWARE 398
5.268. QPID-CPP, PYTHON-QPID, AND SASLWRAPPER 399
5.269. QPID 400
5.270. QT 400
5.271. QUAGGA 402
5.272. QUOTA 403
5.273. RDESKTOP 404
5.274. RDMA 405
5.275. RDMA 406
5.276. READLINE 409
5.277. REDHAT-RELEASE 409
5.278. REDHAT-RPM-CONFIG 410
5.279. RED HAT ENTERPRISE LINUX RELEASE NOTES 411
5.280. RESOURCE-AGENTS 411
5.281. RGMANAGER 413
5.282. RHN-CLIENT-TOOLS AND YUM-RHN-PLUGIN 414
5.283. RICCI 417
5.284. RPCBIND 419
5.285. RPMDEVTOOLS 419
5.286. RPM 420
5.287. RSYNC 421
5.288. RSYSLOG 422
5.289. RUSERS 424
5.290. S390UTILS 424
5.291. SAMBA 427
5.292. SANLOCK 428
5.293. SBLIM-CIM-CLIENT2 428
5.294. SCSI-TARGET-UTILS 429
5.295. SDL 429
5.296. SEABIOS 430
5.297. SED 432
5.298. SELINUX-POLICY 432
5.299. SERVICELOG 443
5.300. SETROUBLESHOOT 444
5.301. SETUP 445
5.302. SLAPI-NIS 447
5.303. SLF4J 447
5.304. SMARTMONTOOLS 448
5.305. SOS 448
5.306. SPICE-CLIENT 451
5.307. SPICE-GTK 452
5.308. SPICE-PROTOCOL 454
5.309. SPICE-SERVER 454
5.310. SPICE-XPI 455
5.311. SQUID 457
5.312. SSSD 457
5.313. STRACE 465
5.314. SUBSCRIPTION-MANAGER 465
5.315. SUBVERSION AND NEON 469
5.316. SUDO 469
5.317. SYSFSUTILS 472

5.318. SYSLINUX 472
5.319. SYSSTAT 473
5.320. SYSTEM-CONFIG-DATE-DOCS 474
5.321. SYSTEM-CONFIG-KDUMP 474
5.322. SYSTEM-CONFIG-KEYBOARD 476
5.323. SYSTEM-CONFIG-LANGUAGE 477
5.324. SYSTEM-CONFIG-LVM 477
5.325. SYSTEM-CONFIG-PRINTER 478
5.326. SYSTEM-CONFIG-USERS 478
5.327. SYSTEMTAP 479
5.328. TAR 481
5.329. TBOOT 483
5.330. TCPDUMP 483
5.331. TELNET 483
5.332. THUNDERBIRD 484
5.333. TOG-PEGASUS 491
5.334. TOMCAT6 492
5.335. TRACE-CMD 493
5.336. TSCLIENT 494
5.337. TUNED 494
5.338. TZDATA 495
5.339. UDEV 498
5.340. UNIXODBC 499
5.341. UPSTART 500
5.342. USBREDIR 500
5.343. UTIL-LINUX-NG 501
5.344. VALGRIND 502
5.345. VIM 502
5.346. VINO 503
5.347. VIOS-PROXY 504
5.348. VIRTIO-WIN 504
5.349. VIRT-MANAGER 506
5.350. VIRT-TOP AND OCAML-LIBVIRT 508
5.351. VIRT-V2V 508
5.352. VIRT-VIEWER 510
5.353. VIRT-WHO 512
5.354. VSFTPD 513
5.355. WGET 515
5.356. WORDNET 515
5.357. WPA_SUPPLICANT 516
5.358. XFIG 517
5.359. XFSPROGS 517
5.360. XINETD 518
5.361. XMLRPC-C 519
5.362. XORG-X11-DRV-ATI AND MESA 519
5.363. XORG-X11-DRV-INTEL 520
5.364. XORG-X11-DRV-MGA 521
5.365. XORG-X11-DRV-QXL 521
5.366. XORG-X11-DRV-WACOM 522
5.367. XORG-X11-SERVER 522
5.368. XULRUNNER 524
5.369. YABOOT 524
5.370. YUM 525

5.371. YUM-UTILS 526
5.372. ZSH 527
5.373. RHNLIB 528
5.374. RHN-CLIENT-TOOLS 528

APPENDIX A. REVISION HISTORY 530

PREFACE

The *Red Hat Enterprise Linux 6.3 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 6.2 and minor release Red Hat Enterprise Linux 6.3.

For system administrators and others planning Red Hat Enterprise Linux 6.3 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 6.3 Technical Notes* provide a single, organized source for change tracking and compliance testing.

For every user, the *Red Hat Enterprise Linux 6.3 Technical Notes* provide details of what has changed in this new release.



NOTE

The Package Manifest is available as a separate document.

CHAPTER 1. 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 6.3. These changes include added or updated **procfs** entries, **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes. For more details on the features added and bugs fixed in the Red Hat Enterprise Linux 6.3 kernel, refer to the Kernel chapter in the 6.3 Release Notes, or Section 5.135.14, " RHSA-2012:0862 – Moderate: Red Hat Enterprise Linux 6.3 kernel security, bug fix, and enhancement update " in this book.

pci=use_crs

The ***pci=use_crs*** boot parameter no longer needs to be specified to force PCI resource allocations to correspond to a specific host bridge the device resides on. It is now the default behavior.

CONFIG_HPET_MMAP, *hpet_mmap*

The high-resolution timer's capacity to remap the HPET registers into the memory of a user process has been enabled via the **CONFIG_HPET_MMAP** option. Additionally, the ***hpet_mmap*** kernel parameter has been added.

pcie_p=nomsi

The ***pcie_p=nomsi*** kernel parameter has been added to allow users to disable MSI/MSI-X for PCI Express Native Hotplug (that is, the **pciehp** driver). When enabled all PCIe ports use INTx for hotplug services.

msi_irqs

A per-PCI device subdirectory has been added to sysfs: **/sys/bus/pci/devices/*<device>*/msi_irqs**. This subdirectory exports the set of MSI vectors allocated by a given PCI device, by creating a numbered subdirectory for each vector under **msi_irqs**. For each vector, various attributes can be exported. Currently the only attribute, named **mode**, tracks the operational mode of that vector (MSI versus MSI-X).

CONFIG_PCI_DEBUG

When the **CONFIG_PCI_DEBUG=y** option is configured, the **-DDEBUG** flag is automatically added to the **EXTRA_CFLAGS** compilation flags.

CONFIG_STRICT_DEVMEM

The **CONFIG_STRICT_DEVMEM** option is enabled by default for the PowerPC architecture. This option restricts access to the **/dev/mem** device. If this option is disabled, userspace access to all memory is allowed, including kernel and userspace memory, and accidental memory (write) access could potentially be harmful.

kdump/kexec configuration options

The following kernel configuration options were enabled for the kdump/kexec kernel dumping mechanism on IBM System z:

```
CONFIG_KEXEC_AUTO_RESERVE=y
CONFIG_CRASH_DUMP=y
CONFIG_PROC_VMCORE=y
```

KEXEC_AUTO_THRESHOLD

The default value for the **KEXEC_AUTO_THRESHOLD** option has been lowered to 2 GB.

/proc/mounts

The **/proc/mounts** file now shows the following mount options for CIFS under the ***dir_mode=*** parameter:

```
nostrictsync
noperm
backupuid
backupgid
```

dmesg_restrict

Writing to the **/proc/sys/kernel/dmesg_restrict** file is only allowed for a root user that has the **CAP_SYS_ADMIN** identifier set.

printk.always_kmsg_dump

A new kernel parameter, ***printk.always_kmsg_dump***, has been added to save the final kernel messages to the reboot, halt, poweroff, and emergency_restart paths. For usage information, refer to the **/usr/share/doc/kernel-doc-*<version>*/Documentation/kernel-parameters.txt** file.

ulimit

The default hard **ulimit** on the number of files has been increased to **4096**:

```
~]$ ulimit -Hn
4096
```

soft_panic

A watchdog module parameter, ***soft_panic***, has been added. When ***soft_panic*** is set to **1**, it causes softdog to invoke kernel panic instead of a reboot when the softdog timer expires. By invoking kernel panic, the system executes kdump, if kdump is configured. Kdump then generates a vmcore which provides additional information on the reasons of the failure.

perf examples

The **/usr/share/doc/perf-*<version>*/examples.txt** documentation file has been added to the perf package.

shm_rmid_forced

Support for the **shm_rmid_forced** sysctl option has been added. When set to **1**, all shared memory objects not referenced in current ipc namespace (with no tasks attached to it) will be automatically forced to use IPC_RMID. For more information refer to **/usr/share/doc/kernel-doc-*<version>*/Documentation/sysctl/kernel.txt** file.

UV systems reduced boot time

A number of patches have been applied to the kernel in Red Hat Enterprise Linux 6.3 to improve overall performance and reduce boot time on extremely large UV systems (patches were tested on a system with 2048 cores and 16 TB of memory). Additionally, boot messages for the SGI UV2 platform were updated.

accept_local

The **/proc/sys/net/ipv4/conf/*/accept_local** sysctl setting has been added to allow a system to receive packets it sent itself. This is needed in order to work with certain load balancing solutions that load balance to themselves.

CONFIG_VGA_SWITCHEROO

The **CONFIG_VGA_SWITCHEROO** configuration option is now enabled by default to allow switching between two graphics cards.

O_DIRECT in FUSE

Support for the **O_DIRECT** flag for files in FUSE (File system in Userspace) has been added.

CONFIG_IP_MROUTE_MULTIPLE_TABLES

The **CONFIG_IP_MROUTE_MULTIPLE_TABLES=y** has been added to enable support for multiple independent multicast routing instances.

nfs.max_session_slots

The ***nfs.max_session_slots*** module/kernel boot parameter has been added. This parameter sets the maximum number of session slots that an NFS client attempts to negotiate with the server.

Default mount option for /proc

In Red Hat Enterprise Linux 6.3, the default mount option of **/proc** during boot up has been changed to:

```
~]# mount -t proc -o nosuid,noexec,nodev proc /proc
```

For third party modules which create devices via **procfs**, please remount **procfs** with the old option:

```
~]# mount -t proc /proc /proc
```

CHAPTER 2. TECHNOLOGY PREVIEWS

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 to fully support Technology Preview features in a future release.

2.1. STORAGE AND FILE SYSTEMS

LVM support for (non-clustered) thinly-provisioned snapshots

A new implementation of LVM copy-on-write (cow) snapshots is available in Red Hat Enterprise Linux 6.3 as a Technology Preview. The main advantage of this implementation, compared to the previous implementation of snapshots, is that it allows many virtual devices to be stored on the same data volume. This implementation also provides support for arbitrary depth of recursive snapshots (snapshots of snapshots of snapshots ...).

This feature is for use on a single system. It is not available for multi-system access in cluster environments.

For more information, refer to the documentation of the **-s**/**--snapshot** option in the **lvcreate** man page.

Package: lvm2-2.02.95-10

LVM support for (non-clustered) thinly-provisioned LVs

Logical Volumes (LVs) can now be thinly provisioned to manage a storage pool of free space to be allocated to an arbitrary number of devices when needed by applications. This allows creation of devices that can be bound to a thinly provisioned pool for late allocation when an application actually writes to the pool. The thinly-provisioned pool can be expanded dynamically if and when needed for cost-effective allocation of storage space. In Red Hat Enterprise Linux 6.3, this feature is introduced as a Technology Preview. You must have the device-mapper-persistent-data package installed to try out this feature. For more information, refer to the **lvcreate(8)** man page.

Package: lvm2-2.02.95-10

Dynamic aggregation of LVM metadata via lvmetad

Most LVM commands require an accurate view of the LVM metadata stored on the disk devices on the system. With the current LVM design, if this information is not available, LVM must scan all the physical disk devices in the system. This requires a significant amount of I/O operations in systems that have a large number of disks.

The purpose of the **lvmetad** daemon is to eliminate the need for this scanning by dynamically aggregating metadata information each time the status of a device changes. These events are signaled to **lvmetad** by **udev** rules. If **lvmetad** is not running, LVM performs a scan as it normally would.

This feature is provided as a Technology Preview and is disabled by default in Red Hat Enterprise Linux 6.3. To enable it, refer to the ***use_lvmetad*** parameter in the **/etc/lvm/lvm.conf** file, and enable the **lvmetad** daemon by configuring the **lvm2-lvmetad** init script.

Package: lvm2-2.02.95-10

Parallel NFS

Parallel NFS (pNFS) is a part of the NFS v4.1 standard that allows clients to access storage devices directly and in parallel. The pNFS architecture eliminates the scalability and performance issues associated with NFS servers in deployment today.

pNFS supports 3 different storage protocols or layouts: files, objects and blocks. The Red Hat Enterprise Linux 6.3 NFS client supports the files layout protocol.

To automatically enable the pNFS functionality, create the **/etc/modprobe.d/dist-nfsv41.conf** file with the following line and reboot the system:

```
alias nfs-layouttype4-1 nfs_layout_nfsv41_files
```

Now when the **-o minorversion=1** mount option is specified, and the server is pNFS-enabled, the pNFS client code is automatically enabled.

For more information on pNFS, refer to http://www.pnfs.com/.

Package: kernel-2.6.32-279

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

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 6.3, SIGAR is considered a Technology Preview package.

Package: sigar-1.6.5-0.4.git58097d9

fsfreeze

Red Hat Enterprise Linux 6 includes **fsfreeze** as a Technology Preview. **fsfreeze** is a new command that halts access to a file system on a disk. **fsfreeze** is designed to be used with hardware RAID devices, assisting in the creation of volume snapshots. For more details on the **fsfreeze** utility, refer to the **fsfreeze(8)** man page.

Package: util-linux-ng-2.17.2-12.7

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-2.6.32-279

Filesystem in user space

Filesystem in Userspace (FUSE) allows for custom file systems to be developed and run in user space.

Package: fuse-2.8.3-4

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-0.19-12

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: lvm2-2.02.95-4

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

2.2. NETWORKING

QFQ queuing discipline

In Red Hat Enterprise Linux 6.3, 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-2.6.32-279

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-0.1-1

IPv6 support in IPVS

The IPv6 support in IPVS (IP Virtual Server) is considered a Technology Preview.

Package: kernel-2.6.32-279

2.3. CLUSTERING AND HIGH AVAILABILITY

Utilizing CPG API for inter-node locking

Rgmanager includes a feature which enables it to utilize Corosync's Closed Process Group (CPG) API for inter-node locking. This feature is automatically enabled when Corosync's RRP feature is enabled. Corosync's RRP feature is considered fully supported. However, when used with the rest of the High-Availability Add-Ons, it is considered a Technology Preview.

Package: rgmanager-3.0.12.1-12

Support for redundant ring for standalone Corosync, BZ#722469

Red Hat Enterprise Linux 6.3 includes support for redundant ring with autorecovery feature as a Technology Preview. Refer to Section 3.7, “Clustering” for a list of known issues associated with this Technology Preview.

Package: corosync-1.4.1-7

corosync-cpgtool, BZ#688260

The **corosync-cpgtool** now specifies both interfaces in a dual ring configuration. This feature is a Technology Preview.

Package: corosync-1.4.1-7

Disabling rgmanager in /etc/cluster.conf, BZ#723925

As a consequence of converting the **/etc/cluster.conf** configuration file to be used by **pacemaker**, **rgmanager** must be disabled. The risk of not doing this is high; after a successful conversion, it would be possible to start **rgmanager** and **pacemaker** on the same host, managing the same resources.

Consequently, Red Hat Enterprise Linux 6 includes a feature (as a Technology Preview) that forces the following requirements:

- **rgmanager** must refuse to start if it sees the **<rm disabled="1">** flag in **/etc/cluster.conf**.
- **rgmanager** must stop any resources and exit if the **<rm disabled="1">** flag appears in **/etc/cluster.conf** during a reconfiguration.

Package: rgmanager-3.0.12.1-12

libqb package

The libqb package provides 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. This package is introduced as a dependency of the pacemaker package, and is considered a Technology Preview in Red Hat Enterprise Linux 6.3.

Package: libqb-0.9.0-2

pacemaker, BZ#456895

Pacemaker, a scalable high-availability cluster resource manager, is included in Red Hat Enterprise Linux 6 as a Technology Preview. Pacemaker is not fully integrated with the Red Hat cluster stack.

Package: pacemaker-1.1.7-6

2.4. AUTHENTICATION

Support for central management of SSH keys, BZ#803822

Previously, it was not possible to centrally manage host and user SSH public keys. Red Hat Enterprise Linux 6.3 includes SSH public key management for Identity Management servers as a Technology Preview. OpenSSH on Identity Management clients is automatically configured to use public keys which are stored on the Identity Management server. SSH host and user identities can now be managed centrally in Identity Management.

Package: sssd-1.8.0-32

SELinux user mapping, BZ#803821

Red Hat Enterprise Linux 6.3 introduces the ability to control the SELinux context of a user on a remote system. SELinux user map rules can be defined and, optionally, associated with HBAC rules. These maps define the context a user receives depending on the host they are logging into and the group membership. When a user logs into a remote host which is configured to use SSSD with the Identity Management backend, the user's SELinux context is automatically set according to mapping rules defined for that user. For more information, refer to http://freeipa.org/page/SELinux_user_mapping. This feature is considered a Technology Preview.

Package: sssd-1.8.0-32

SSSD support for automount map caching, BZ#761570

In Red Hat Enterprise Linux 6.3, SSSD includes a new Technology Preview feature: support for caching automount maps. This feature provides several advantages to environments that operate with **autofs**:

- Cached automount maps make it easy for a client machine to perform mount operations even when the LDAP server is unreachable, but the NFS server remains reachable.
- When the **autofs** daemon is configured to look up automount maps via SSSD, only a single file has to be configured: **/etc/sssd/sssd.conf**. Previously, the **/etc/sysconfig/autofs** file had to be configured to fetch autofs data.
- Caching the automount maps results in faster performance on the client and lower traffic on the LDAP server.

Package: sssd-1.8.0-32

2.5. 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 in Red Hat Enterprise Linux 6.3.

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

2.6. DEVICES

SR-IOV on the be2net driver, BZ#602451

The SR-IOV functionality of the Emulex **be2net** driver is considered a Technology Preview in Red Hat Enterprise Linux 6.3. You must meet the following requirements to use the latest version of SR-IOV support:

- You must run the latest Emulex firmware (revision 4.1.417.0 or later).
- The server system BIOS must support the SR-IOV functionality and have virtualization support for Direct I/O VT-d.
- You must use the GA version of Red Hat Enterprise Linux 6.3.

SR-IOV runs on all Emulex-branded and OEM variants of BE3-based hardware, which all require the **be2net** driver software.

Package: kernel-2.6.32-279

iSCSI and FCoE boot

iSCSI and FCoE boot support on Broadcom devices is not included in Red Hat Enterprise Linux 6.3. These two features, which are provided by the **bnx2i** and **bnx2fc** Broadcom drivers, remain a Technology Preview until further notice.

Package: kernel-2.6.32-279

mpt2sas lockless mode

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

Package: kernel-2.6.32-279

2.7. KERNEL

Thin-provisioning and scalable snapshot capabilities

The **dm-thinp** targets, **thin** and **thin-pool**, provide a device mapper device with thin-provisioning and scalable snapshot capabilities. This feature is available as a Technology Preview.

Package: kernel-2.6.32-279

kdump/kexec kernel dumping mechanism for IBM System z

In Red Hat Enterprise Linux 6.3, the kdump/kexec kernel dumping mechanism is enabled for IBM System z systems as a Technology Preview, in addition to the IBM System z stand-alone and hypervisor dumping mechanism. The auto-reserve threshold is set at 4 GB; therefore, any IBM System z system with more than 4 GB of memory has the kexec/kdump mechanism enabled.

Sufficient memory must be available because kdump reserves approximately 128 MB as default. This is especially important when performing an upgrade to Red Hat Enterprise Linux 6.3. Sufficient disk space must also be available for storing the dump in case of a system crash. Kdump is limited to DASD or QETH networks as dump devices until kdump on SCSI disk is supported.

The following warning message may appear when kdump is initialized:

```
..no such file or directory
```

This message does not impact the dump functionality and can be ignored. You can configure or disable kdump via **/etc/kdump.conf**, **system-config-kdump**, or **firstboot**.

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-2.6.32-279

Remote audit logging

The audit package contains the user space utilities for storing and searching the audit records generated by the **audit** subsystem in the Linux 2.6 kernel. Within the audispd-plugins sub-package is a utility that allows for the transmission of audit events to a remote aggregating machine. This remote audit logging application, **audisp-remote**, is considered a Technology Preview in Red Hat Enterprise Linux 6.

Package: audispd-plugins-2.2-2

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-3.1.5-17

2.8. VIRTUALIZATION

Performance monitoring in KVM guests, BZ#645365

KVM can now virtualize a performance monitoring unit (vPMU) to allow virtual machines to use performance monitoring. Additionally it supports Intel's "architectural PMU" which can be live-migrated across different host CPU versions, using the **-cpu** host flag.

With this feature, Red Hat virtualization customers are now able to utilize performance monitoring in KVM guests seamlessly. The virtual performance monitoring feature allows virtual machine users to identify sources of performance problems in their guests, using their preferred pre-existing profiling tools that work on the host as well as the guest. This is an addition to the existing ability to profile a KVM guest from the host.

This feature is a Technology Preview in Red Hat Enterprise Linux 6.3.

Package: kernel-2.6.32-279

Dynamic virtual CPU allocation

KVM in Red Hat Enterprise Linux 6.3 now supports dynamic virtual CPU allocation, also called vCPU hot plug, to dynamically manage capacity and react to unexpected load increases on their platforms during off-peak hours.

The virtual CPU hot-plugging feature gives system administrators the ability to dynamically adjust CPU resources in a guest. Because a guest no longer has to be taken offline to adjust the CPU resources, the availability of the guest is increased.

This feature is a Technology Preview in Red Hat Enterprise Linux 6.3. Currently, only the vCPU hot-add functionality works. The vCPU hot-unplug feature is not yet implemented.

Package: qemu-kvm-0.12.1.2-2.295

Virtio-SCSI capabilities

KVM Virtualization's storage stack has been improved with the addition of virtio-SCSI (a storage

architecture for KVM based on SCSI) capabilities. Virtio-SCSI provides the ability to connect directly to SCSI LUNs and significantly improves scalability compared to virtio-blk. The advantage of virtio-SCSI is that it is capable of handling hundreds of devices compared to virtio-blk which can only handle 25 devices and exhausts PCI slots.

Virtio-SCSI is now capable of inheriting the feature set of the target device with the ability to:

- attach a virtual hard drive or CD through the virtio-scsi controller,
- pass-through a physical SCSI device from the host to the guest via the QEMU scsi-block device,
- and allow the usage of hundreds of devices per guest; an improvement from the 32-device limit of virtio-blk.

This feature is a Technology Preview in Red Hat Enterprise Linux 6.3

Package: qemu-kvm-0.12.1.2-2.295

Support for in-guest S4/S3 states

KVM's power management features have been extended to include native support for S4 (suspend to disk) and S3 (suspend to RAM) states within the virtual machine, speeding up guest restoration from one of these low power states. In earlier implementations guests were saved or restored to/from a disk or memory that was external to the guest, which introduced latency.

Additionally, machines can be awakened from S3 with events from a remote keyboard through SPICE.

This feature is a Technology Preview and is disabled by default in Red Hat Enterprise Linux 6.3. To enable it, select the **/usr/share/seabios/bios-pm.bin** file for the VM bios instead of the default **/usr/share/seabios/bios.bin** file.

The native, in-guest S4 (suspend to disk) and S3 (suspend to RAM) power management features support the ability to perform suspend to disk and suspend to RAM functions in the guest (as opposed to the host), reducing the time needed to restore a guest by responding to simple keyboard gestures input. This also removes the need to maintain an external memory-state file. This capability is supported on Red Hat Enterprise Linux 6.3 guests and Windows guests running on any hypervisor capable of supporting S3 and S4.

Package: seabios-0.6.1.2-19

System monitoring via SNMP, BZ#642556

This feature provides KVM support for stable technology that is already used in data center with bare metal systems. SNMP is the standard for monitoring and is extremely well understood as well as computationally efficient. System monitoring via SNMP in Red Hat Enterprise Linux 6 allows the KVM hosts to send SNMP traps on events so that hypervisor events can be communicated to the user via standard SNMP protocol. This feature is provided through the addition of a new package: libvirt-snmp. This feature is introduced as a Technology Preview.

Package: libvirt-snmp-0.0.2-3

Wire speed requirement in KVM network drivers

Virtualization and cloud products that run networking work loads need to run wire speeds. Up until Red Hat Enterprise Linux 6.1, the only way to reach wire speed on a 10 GB Ethernet NIC with a lower CPU utilization was to use PCI device assignment (passthrough), which limits other features like

memory overcommit and guest migration

The **macvtap/vhost** zero-copy capabilities allow the user to use those features when high performance is required. This feature improves performance for any Red Hat Enterprise Linux 6.x guest in the VEPA use case. This feature is introduced as a Technology Preview.

Package: qemu-kvm-0.12.1.2-2.295

2.9. RESOURCE MANAGEMENT

numad package

The numad package provides a daemon for NUMA (Non-Uniform Memory Architecture) systems that monitors NUMA characteristics. As an alternative to manual static CPU pining and memory assignment, numad provides dynamic adjustment to minimize memory latency on an ongoing basis. The package also provides an interface that can be used to query the **numad** daemon for the best manual placement of an application. The numad package is introduced as a Technology Preview.

Package: numad-0.5-4.20120522git

CHAPTER 3. KNOWN ISSUES

3.1. INSTALLATION

anaconda component

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

anaconda component

Installing Red Hat Enterprise Linux 6.3 using the text user interface on a system which already has a Red Hat Enterprise Linux system installed on the disk, and going back to the initial Anaconda installation page (using the Back button) may cause a traceback error.

dracut component

Installations to a network root device, such as an iSCSI device, do not function properly when using DHCP, preventing the installed system from rebooting. To work around this issue, when installing to an iSCSI root device, you must select the Anaconda installer option **Bind targets to network interfaces**; do not leave it unselected, as is the default. Additionally, you must use static IP addresses if using a network root device.

To work around this issue when installing via kickstart, add the **--iface=** option to the iSCSI command, for example:

```
iscsi --ipaddr 10.34.39.46 --port 3260 --target iqn.2009-02.com.kvm:iscsibind --iface=eth0
```

anaconda component

Red Hat Enterprise Linux 6.3 fails to boot when installed without LVM and booted from a Storage Area Network (SAN). To work around this issue, ensure that the **/boot** partition is using the first partition of multipath, or use LVM (which is the default behavior).

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 **clearpart** 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

In Red Hat Enterprise Linux 6.3, Anaconda allows installation to disks of size 2.2 TB and larger, but the installed system may not boot properly. Disks of size 2.2 TB and larger may be used during the installation process, but only as data disks (that is, should not be used as bootable disks).

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.

squashfs-tools component

During the installation on POWER systems, error messages similar to the following may be returned to sys.log:

```
attempt to access beyond end of device
loop0: rw=0, want=248626, limit=248624
```

These errors do not prevent installation and only occur during the initial setup. The file system created by the installer will function 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.

3.2. ENTITLEMENT

subscription manager component

When registering a system with **firstboot**, the *RHN Classic* option is checked by default in the Subscription part.

subscription manager component, BZ#811771

Subscription Manager now disables **gpgcheck** for any repositories it manages which have an empty **gpgkey**. To re-enable the repository, upload the GPG keys, and ensure that the correct URL is added to your custom content definition.

3.3. DEPLOYMENT

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.

3.4. VIRTUALIZATION

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.

vdsm component, BZ#826921

The following parameter has been deprecated in the **/etc/vdsm/vdsm.conf** file:

```
[irs]
nfs_mount_options = soft,nosharecache,vers=3
```

This parameter will continue to be supported in versions 3.x, but will be removed in version 4.0 of NFS. Customers using this parameter should upgrade their domains to V2 and greater and set the parameters from the GUI.

vdsm component, BZ#749479

When adding a bond to an existing network, its world-visible MAC address may change. If the DHCP server is not aware that the new MAC address belongs to the same host as the old one, it may assign the host a different IP address, that is unknown to the DNS server nor to Red Hat Enterprise Virtualization Manager. As a result, Red Hat Enterprise Virtualization Manager VDSM connectivity is broken.

To work around this issue, configure your DHCP server to assign the same IP for all the MAC addresses of slave NICs. Alternatively, when editing a management network, do not check connectivity, and make sure that Red Hat Enterprise Virtualization Manager / DNS use the newly-assigned IP address for the node.

vdsm component

Vdsm uses cgroups if they are available on the host. If the **cgconfig** service is turned off, turn it on with the **chkconfig cgconfig on** command and reboot. If you prefer not to reboot your system, restarting the **libvirt** service and **vdsm** should be sufficient.

ovirt-node component, BZ#747102

Upgrades from Beta to the GA version will result in an incorrect partitioning of the host. The GA version must be installed clean. UEFI machines must be set to legacy boot options for RHEV-H to boot successfully after installation.

kernel component

When a system boots from SAN, it starts the **libvirtd** service, which enables IP forwarding. The service causes a driver reset on both Ethernet ports which causes a loss of all paths to an OS disk. Under this condition, the system cannot load firmware files from the OS disk to initialize Ethernet ports, eventually never recovers paths to the OS disk, and fails to boot from SAN. To work around this issue add the **bnx2x.disable_tpa=1** option to the kernel command line of the GRUB menu, or do not install virtualization related software and manually enable IP forwarding when needed.

vdsm component

If the **/root/.ssh/** directory is missing from a host when it is added to a Red Hat Enterprise Virtualization Manager data center, the directory is created with a wrong SELinux context, and SSH'ing into the host is denied. To work around this issue, manually create the **/root/.ssh** directory with the correct SELinux context:

```
~]# mkdir /root/.ssh
~]# chmod 0700 /root/.ssh
~]# restorecon /root/.ssh
```

vdsm component

VDSM now configures **libvirt** so that connection to its local read-write UNIX domain socket is password-protected by SASL. The intention is to protect virtual machines from human errors of local host administrators. All operations that may change the state of virtual machines on a Red Hat Enterprise Virtualization-controlled host must be performed from Red Hat Enterprise Virtualization Manager.

libvirt component

In earlier versions of Red Hat Enterprise Linux, **libvirt** permitted PCI devices to be insecurely assigned to guests. In Red Hat Enterprise Linux 6, assignment of insecure devices is disabled by default by **libvirt**. However, this may cause assignment of previously working devices to start failing. To enable the old, insecure setting, edit the **/etc/libvirt/qemu.conf** file, set the ***relaxed_acs_check = 1*** parameter, and restart **libvirtd** (**service libvirtd restart**). Note that this action will re-open possible security issues.

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 often crashes during the final steps of the installation process due to a system hang. 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.

3.5. STORAGE AND FILE SYSTEMS

Driver Update Disk component

The hpsa driver installed from the AMD64 and Intel 64 Driver Update Program ISO might not be loaded properly on Red Hat Enterprise Linux 6.3. Consequently, the system can become unresponsive. To work around this problem, use the **pci=nomsi** kernel parameter before installing the driver from the ISO.

lvm2 component, BZ#832392

When **issue_discards=1** is configured in the **/etc/lvm/lvm.conf** file, moving physical volumes via the **pvmove** command results in data loss. To work around this issue, ensure that **issue_discards=0** is set in your **lvm.conf** file before moving any physical volumes.

lvm2 component, BZ#832033

When using the **lvmetad** daemon (currently a Technology Preview), avoid passing the **--test** argument to commands. The use of the **--test** argument may lead to inconsistencies in the cache that **lvmetad** maintains. This issue will be fixed in a future release. If the **--test** argument has been used, fix any problems by restarting the **lvmetad** daemon.

lvm2 component, BZ#820229

It is not possible to rename thin logical volumes using tools provided in the current LVM2 release. The rename operation returns the following error:

```
lvrename Cannot rename <volume_name>: name format not recognized for internal LV
<pool_name>
```

This issue will be fixed in the next LVM2 release.

device-mapper-multipath component

Multipath's **queue_without_daemon yes** default option queues I/O even though all iSCSI links have been disconnected when the system is shut down, which causes LVM to become unresponsive when scanning all block devices. As a result, the system cannot be shut down. To work around this issue, add the following line into the **defaults** section of **/etc/multipath.conf**:

```
queue_without_daemon no
```

initscripts component

Running the file system check (using **fsck**) on a NFS mounted file system fails, and causes the system to fail to boot and drop into a shell. To work around this issue, disable **fsck** on any **/boot** partitions by setting the sixth value of a **/boot** entry in **/etc/fstab** to **0**.

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 **dracut** utility currently only supports one FiberChannel over Ethernet (FCoE) connection to be used to boot from the root device. Consequently, booting from a root device that spans multiple FCoE devices (for example, using RAID, LVM or similar techniques) is not possible.

***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>
```

3.6. NETWORKING

kernel component

Some e1000e 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-eth*<X>*** 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 6.3 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, BZ#816888

Running the QFQ queuing discipline in a virtual guest eventually results in kernel panic.

kernel component

Because RHEL6.3 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.

perftest component

The **rdma_bw** and **rdma_lat** utilities (provided by the perftest package) are now deprecated and will be removed from the perftest package in a future update. Users should use the following utilities instead: **ib_write_bw**, **ib_write_lat**, **ib_read_bw**, and **ib_read_lat**.

3.7. CLUSTERING

corosync component, BZ#722469

A double ring failure results in the spinning of the corosync process. Also, because DLM relies on SCTP, which is non-functional, many features of the cluster software that rely on DLM do not work properly.

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.

3.8. AUTHENTICATION

Identity Management component

When using the Identity Management WebUI in the Internet Explorer browser, you may encounter the following issues:

- While the browser window is not maximized or many users are logged into the WebUI, scrolling down a page to select a user may not work properly. As soon as the user's checkbox is selected, the scroll bar jumps back up without selecting the user. This error also occurs when a permission is added to a privilege. (BZ#831299)

- When attempting to edit a service, the edit page for that service may occasionally be blank, or show only labels for **Principal** or **Service** without showing their values. When adding a service, under certain conditions, the drop-down menu lists the available services and hosts but users are unable to select any of the entries. (BZ#831227)

- When adding a permission of type subtree, the text area to specify the subtree is too small and non-resizable making it difficult to enter long subtree entries. (BZ#830817)

- When adding a delegation, its attributes are separated by disproportionately large vertical spaces. (BZ#829899)

- When adding a member, the edge of the displayed window suggests it can be resized. However, resizing of the window does not work. When adding a Sudo Command to a Sudo Command group, the first group overlays with the column title. (BZ#829746)

- Adding a new DNS zone causes the window to be incorrectly rendered as text on the existing page. (BZ#827583)

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 come 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#783502

The Identity Management permission plug-in does not verify that the set of attributes specified for a new permission is relevant to the target object type that the permission allows access to. This means a user is able to create a permission which allows access to attributes that will never be present in the target object type because such attributes are not allowed in its object classes. You must ensure that the chosen set of attributes for which a new permission grants access to is relevant to the chosen target object type.

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 do 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

Identity Management and the **mod_ssl** module should not be installed on the same system, otherwise Identity Management is unable to issue certificates because **mod_ssl** holds the **mod_proxy** hooks. To work around this issue, uninstall **mod_ssl**.

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, BZ#750922

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, BZ#750596

Two Identity Management servers, both with a CA (Certificate Authority) installed, use two replication 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.

Identity Management component

On the configuration page of the Identity Management WebUI, if the **User** search field is left blank, and the **search** button is clicked, an internal error is returned.

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/sssd/sssd.conf** file:

```
ldap_referrals = false
```

3.9. DEVICES

ipmitool component

Not specifying the **-N** option when setting retransmission intervals of IPMI messages over the LAN or LANplus interface may cause various error messages to be returned. For example:

```
~]# ipmitool -I lanplus -H $HOST -U root -P $PASS sensor list
Unable to renew SDR reservation
Close Session command failed: Reservation cancelled or invalid

~]# ipmitool -I lanplus -H $HOST -U root -P $PASS delloem powermonitor
Error getting power management information, return code c1
Close Session command failed: Invalid command
```

ipmitool component

The **ipmitool** may crash in certain cases. For example, when an incorrect password is used, a segmentation fault occurs:

```
~]# ipmitool -I lanplus -H $HOST -U root -P wrongpass delloem powermonitor
Error: Unable to establish IPMI v2 / RMCP+ session
Segmentation fault (core dumped)
```

kernel component,

Unloading the **be2net** driver with a Virtual Function (VF) attached to a virtual guest results in kernel panic.

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.

kernel component

iSCSI and FCoE boot support on Broadcom devices is not included in Red Hat Enterprise Linux 6.3. These two features, which are provided by the **bnx2i** and **bnx2fc** Broadcom drivers, remain a Technology Preview until further notice.

kexec-tools component

Starting with Red Hat Enterprise Linux 6.0 and later, kexec kdump supports dumping core to the Brtfs 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.

busybox component

When running kdump in a busybox environment and dumping to a Btrfs file system, you may receive the following error message:

```
/etc/kdump.conf: Unsupported type btrfs
```

However, Btrfs is supported as a kdump target. To work around this issue, install the btrfs-progs package, verify that the **/sbin/btrfsck** file exists, and retry.

trace-cmd component

The **trace-cmd** service does 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.

tuned component

Red Hat Enterprise Linux 6.1 and later enter processor power-saving states more aggressively. This may result in a small performance penalty on certain workloads. This functionality may be disabled at boot time by passing the ***intel_idle.max_cstate=0*** parameter, or at run time by using the **cpu_dma_latency pm_qos** interface.

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.

3.10. KERNEL

kernel component

Intel Xeon E5-XXXX V2 Series Processor running on the C600 chipset is not supported in Red Hat Enterprise Linux 6.3. An "unsupported hardware" message can therefore be reported by the kernel.

kernel component

The Red Hat Enterprise Linux 6.3 kernels upgraded the **mlx4** modules to a later version. If the modules are used together with, for example, the HP InfiniBand Enablement Kit, the behavior is different. Consequently, certain Mellanox cards do not come up with network interfaces on Red Hat Enterprise Linux 6.3. To work around this problem, the **mlx7_core** module has to be loaded with the **port_type_array** option and a **2** parameter for each used InfiniBand card. Follow this example to manually load the driver for two cards in the system:

```
~]# rmmod mlx4_en
~]# rmmod mlx4_core
~]# modprobe mlx4_core port_type_array=2,2
~]# modprobe mlx4_en
~]# ip a
```

The last of the above commands will show the new interfaces. To configure these parameters to be applied by the system when the modules are loaded, run:

```
~]# echo 'options mlx4_core port_type_array=2,2' >/etc/modprobe.d/mlx4_core.conf
```

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

In Red Hat Enterprise Linux 6.3, three module parameters (***num_lro***, ***rss_mask***, and ***rss_xor***) that were supported by older versions of the **mlx4_en** driver have become obsolete and are no longer used. If you supply these parameters, the Red Hat Enterprise Linux 6.3 driver will ignore them and log a warning.

kernel component

Due to a race condition, in certain cases, writes to RAID4/5/6 while the array is reconstructing could hang the system.

kernel component

The installation of Red Hat Enterprise Linux 6.3 i386 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

Platforms with BIOS/UEFI that are unaware of PCI-e SR-IOV capabilities may fail to enable virtual functions

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#690523

If appropriate SCSI device handlers (**scsi_dh** modules) are not available when the storage driver (for example, **lpfc**) is first loaded, I/O operations may be issued to SCSI multipath devices that are not ready for those I/O operations. This can result in significant delays during system boot and excessive

I/O error messages in the kernel log.

Provided the storage driver is loaded before **multipathd** is started (which is the default behavior), users can work around this issue by making sure the appropriate SCSI device handlers (**scsi_dh** modules) are available by specifying one of the following kernel command line parameters which dracut consumes:

- rdloaddriver=scsi_dh_emc
- rdloaddriver=scsi_dh_rdac,scsi_dh_hp_sw
- rdloaddriver=scsi_dh_emc,scsi_dh_rdac,scsi_dh_alua

Note that the order of the listed **scsi_dh** modules does not matter.

Specifying one of the above parameters causes the **scsi_dh** module(s) to load before the storage driver is loaded or multipath is started.

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, BZ#770545

In Red Hat Enterprise Linux 6.2 and Red Hat Enterprise Linux 6.3, the default value for **sysctl vm.zone_reclaim_mode** is now **0**, whereas in Red Hat Enterprise Linux 6.1 it was **1**.

kernel component

Using Alsa with an HDA Intel sound card and the Conexant CX20585 codec causes sound and recording failures. To work around this issue, add the following line to the **/etc/modprobe.d/dist-alsa.conf** file:

```
options snd-hda-intel model=thinkpad
```

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

The **lpfc** driver is deprecating the **sysfs mbox** interface as it is no longer used by the Emulex tools. Reads and writes are now stubbed out and only return the **-EPERM** (Operation not permitted) symbol.

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

Enabling CHAP (Challenge-Handshake Authentication Protocol) on an iSCSI target for the **be2iscsi** driver results in kernel panic. To work around this issue, disable CHAP on the iSCSI target.

kernel component

Newer VPD (Vital Product Data) blocks can exceed the size the **tg3** driver normally handles. As a result, some of the routines that operate on the VPD blocks may fail. For example, the **nvram** test fails when running the **ethtool –t** command on BCM5719 and BCM5720 Ethernet Controllers.

kernel component

Running the **ethtool -t** command on BCM5720 Ethernet controllers causes a loopback test failure because the **tg3** driver does not wait long enough for a link.

kernel component

The **tg3** driver in Red Hat Enterprise Linux 6.2 does not include support for Jumbo frames and TSO (TCP Segmentation Offloading) on BCM5719 Ethernet controllers. As a result, the following error message is returned when attempting to configure, for example, Jumbo frames:

```
SIOCSIFMTU: Invalid argument
```

kernel component

The default interrupt configuration for the Emulex LPFC FC/FCoE driver has changed from INT-X to MSI-X. This is reflected by the ***lpfc_use_msi*** module parameter (in **/sys/class/scsi_host/host#/lpfc_use_msi**) being set to **2** by default, instead of the previous **0**.

Two issues provide motivation for this change: SR-IOV capability only works with the MSI-X interrupt mode, and certain recent platforms only support MSI or MSI-X.

However, the change to the LPFC default interrupt mode can bring out host problems where MSI/MSI-X support is not fully functional. Other host problems can exist when running in the INT-X mode.

If any of the following symptoms occur after upgrading to, or installing Red Hat Enterprise Linux 6.2 with an Emulex LPFC adapter in the system, change the value of the **lpfc** module parameter, ***lpfc_use_msi***, to **0**:

- The initialization or attachment of the **lpfc** adapter may fail with mailbox errors. As a result, the **lpfc** adapter is not configured on the system. The following message appear in **/var/log/messages**:

  ```
  lpfc 0000:04:08.0: 0:0:0443 Adapter failed to set maximum DMA length mbxStatus x0
  lpfc 0000:04:08.0: 0:0446 Adapter failed to init (255), mbxCmd x9 CFG_RING, mbxStatus
  x0, ring 0
  ```

```
lpfc 0000:04:08.0: 0:1477 Failed to set up hba
ACPI: PCI interrupt for device 0000:04:08.0 disabled
```

- While the **lpfc** adapter is operating, it may fail with mailbox errors, resulting in the inability to access certain devices. The following message appear in **/var/log/messages**:

```
lpfc 0000:0d:00.0: 0:0310 Mailbox command x5 timeout Data: x0 x700 xffff81039ddd0a00
lpfc 0000:0d:00.0: 0:0345 Resetting board due to mailbox timeout
lpfc 0000:0d:00.0: 0:(0):2530 Mailbox command x23 cannot issue Data: xd00 x2
```

- Performing a warm reboot causes any subsequent boots to halt or stop because the BIOS is detecting the **lpfc** adapter. The system BIOS logs the following messages:

```
Installing Emulex BIOS .....
Bringing the Link up, Please wait...
Bringing the Link up, Please wait...
```

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, BZ#683012

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=lapic*** 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.

3.11. DESKTOP

libwacom component

The Lenovo X220 Tablet Touchscreen is not supported in the kernel shipped with Red Hat Enterprise Linux 6.3.

wacomcpl package, 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.

gnome-settings-daemon component, BZ#826128

On some tablets, using the NVIDIA Graphics drivers to configure Twinview causes the tablet motions to be incorrectly mapped to the laptop screen instead of the tablet itself. Using the stylus on the tablet moves the cursor on the laptop screen.

acroread component

Running a AMD64 system without the sssd-client.i686 package installed, which uses SSSD for getting information about users, causes **acroread** 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.

3.12. TOOLS

matahari component

The Matahari agent framework (matahari-*) packages are deprecated starting with the Red Hat Enterprise Linux 6.3 release. 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. It is strongly recommended that users discontinue the use of the matahari packages and other packages which depend on the Matahari infrastructure (specifically, libvirt-qmf and fence-virtd-libvirt-qpid). It is recommended that users uninstall Matahari from their systems to remove any possibility of security issues being exposed.

Users who choose to continue to use the Matahari agents should note the following:

- The matahari packages are not installed by default starting with Red Hat Enterprise Linux 6.3 and are not enabled by default to start on boot when they are installed. Manual action is needed to both install and enable the **matahari** services.
- The default configuration for **qpid** (the transport agent used by Matahari) does not enable access control lists (ACLs) or SSL. Without ACLs/SSL, the Matahari infrastructure is not secure. Configuring Matahari without ACLs/SSL is not recommended and may reduce your system's security.
- The **matahari-services** agent is specifically designed to allow remote manipulation of services (start, stop). Granting a user access to Matahari services is equivalent to providing a remote user with root access. Using Matahari agents should be treated as equivalent to providing remote root SSH access to a host.
- By default in Red Hat Enterprise Linux, the Matahari broker (**qpidd** running on port **49000**) does not require authentication. However, the Matahari broker is not remotely accessible unless the firewall is disabled, or a rule is added to make it accessible. Given the capabilities exposed by Matahari agents, if Matahari is enabled, system administrators should be extremely cautious with the options that affect remote access to Matahari.

Note that Matahari will not be shipped in future releases of Red Hat Enterprise Linux (including Red Hat Enterprise Linux 7), and may be considered for formal removal in a future release of Red Hat Enterprise Linux 6.

libreport component

An error in the default **libreport** configuration causes the following warning message to appear during problem reporting:

```
/bin/sh: line 4: reporter-bugzilla: command not found
```

This warning message has no effect on the functionality of **libreport**. To prevent the warning message from being displayed, replace the following lines in the **/etc/libreport/events.d/ccpp_event.conf** file:

```
abrt-action-analyze-backtrace &&
(
    bug_id=$(reporter-bugzilla -h `cat duphash`) &&
    if test -n "$bug_id"; then
        abrt-bodhi -r -b $bug_id
    fi
)
```

with:

-

```
abrt-action-analyze-backtrace
```

irqbalance component, BZ#813078

The **irqbalance(1)** man page does not contain documentation for the **IRQBALANCE_BANNED_CPUS** and **IRQBALANCE_BANNED_INTERRUPTS** environment variables. The following documentation will be added to this man page in a future release:

IRQBALANCE_BANNED_CPUS

> Provides a mask of cpus which irqbalance should ignore and never assign interrupts to. This is a hex mask without the leading '0x', on systems with large numbers of processors each group of eight hex digits is sepearated ba a comma ','. i.e. `export IRQBALANCE_BANNED_CPUS=fc0` would prevent irqbalance from assigning irqs to the 7th-12th cpus (cpu6-cpu11) or `export IRQBALANCE_BANNED_CPUS=ff000000,00000001` would prevent irqbalance from assigning irqs to the 1st (cpu0) and 57th-64th cpus (cpu56-cpu63).

IRQBALANCE_BANNED_INTERRUPTS

> Space seperated list of integer irq's which irqbalance should ignore and never change the affinity of. i.e.
>
> export IRQBALANCE_BANNED_INTERRUPTS="205 217 225"

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
```

parted component

The **parted** utility in Red Hat Enterprise Linux 6 cannot handle Extended Address Volumes (EAV) Direct Access Storage Devices (DASD) that have more than 65535 cylinders. Consequently, EAV DASD drives cannot be partitioned using **parted**, and installation on EAV DASD drives will fail. To work around this issue, complete the installation on a non EAV DASD drive, then add the EAV device after the installation using the tools provided in the s390-utils package.

CHAPTER 4. NEW PACKAGES

4.1. RHEA-2012:0842 – NEW PACKAGE: BYZANZ

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

The byzanz package contains an easy-to-use desktop recorder that can record to GIF images, Ogg Theora video (optionally with sound), and other formats. A GNOME panel applet and a command-line recording tool are also included in the package.

This enhancement update adds the byzanz package to Red Hat Enterprise Linux 6. (BZ#623262)

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

4.2. RHEA-2012:0797 – NEW PACKAGES: CRASH-GCORE-COMMAND

New crash-gcore-command packages are now available for Red Hat Enterprise Linux 6.

The crash-gcore-command extension module is used to dynamically add a gcore command to a running crash utility session on a kernel dumpfile. The command will create a core dump file for a specified user task program that was running when a kernel crashed. The resultant core dump file may then be used with gdb.

This enhancement update adds the crash-gcore-command packages to Red Hat Enterprise Linux 6. (BZ#692799)

All users who require the crash-gcore-command should install these new packages.

4.3. RHEA-2012:0831 – NEW PACKAGE: DEVICE-MAPPER-PERSISTENT-DATA

A new device-mapper-persistent-data package is now available for Red Hat Enterprise Linux 6.

The device-mapper-persistent-data package provides device-mapper thin provisioning (thinp) tools.

This enhancement update adds the device-mapper-persistent-data package to Red Hat Enterprise Linux 6 as a Technology Preview. (BZ#760614)

More information about Red Hat Technology Previews is available here:

https://access.redhat.com/support/offerings/techpreview/

All users who require device-mapper-persistent-data should install this new package, which adds this enhancement.

4.4. RHEA-2012:0814 – NEW PACKAGE: I2C-TOOLS

A new i2c-tools package is now available for Red Hat Enterprise Linux 6.

The i2c-tools package contains a set of I2C tools for Linux: a bus probing tool, a chip dumper, register-level SMBus access helpers, EEPROM (Electrically Erasable Programmable Read-Only Memory) decoding scripts, EEPROM programming tools, and a python module for SMBus access.



NOTE

EEPROM decoding scripts can render your system unusable. Make sure to use these tools wisely.

This enhancement update adds the i2c-tools package to Red Hat Enterprise Linux 6. (BZ#773267)

All users who require i2c-tools should install this new package.

4.5. RHEA-2012:0829 – NEW PACKAGES: IPSET AND LIBMNL

New ipset and libmnl packages are now available for Red Hat Enterprise Linux 6.

The ipset packages provide IP sets, a framework inside the Linux 2.4.x and 2.6.x kernel, which can be administered by the ipset utility. Depending on the type, an IP set can currently store IP addresses, TCP/UDP port numbers or IP addresses with MAC addresses in a way that ensures high speed when matching an entry against a set.

The libmnl packages required by the ipset packages provide a minimalistic user-space library oriented to Netlink developers. The library provides functions to make socket handling, message building, validating, parsing, and sequence tracking easier.

This enhancement update adds the ipset and libmnl packages to Red Hat Enterprise Linux 6. (BZ#477115, BZ#789346)

All users who require ipset and libmnl are advised to install these new packages.

4.6. RHEA-2012:0840 – NEW PACKAGES: JAVA-1.7.0-IBM

New java-1.7.0-ibm packages are now available for Red Hat Enterprise Linux 6.

The java-1.7.0-ibm packages provide the IBM Java 7 Runtime Environment and the IBM Java 7 Software Development Kit.

This update adds the java-1.7.0-ibm packages to Red Hat Enterprise Linux 6. (BZ#693783)

Note: Before applying this update, make sure that any previous IBM Java packages have been removed.

All users who require java-1.7.0-ibm should install these new packages.

4.7. RHEA-2012:0981 – NEW PACKAGES: JAVA-1.7.0-OPENJDK

New java-1.7.0-openjdk packages that provide OpenJDK 7 are now available as a Technology Preview for Red Hat Enterprise Linux 6.

[Updated 9 June 2012]

> This advisory has been updated to reflect the fact that java-1.7.0-openjdk is fully supported and no longer claims that java-1.7.0-openjdk is a Technology Preview feature. The packages included in this revised update have not been changed in any way from the packages included in the previous version of this advisory.

The java-1.7.0-openjdk packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Java Software Development Kit.

This enhancement update adds new java-1.7.0-openjdk package to Red Hat Enterprise Linux 6. (BZ#803726)

These packages do not replace the previous version of the OpenJDK (java-1.6.0-openjdk) if present. Users can safely install OpenJDK 7 in addition to OpenJDK 6. The system default version of Java can be configured using the 'alternatives' tool.

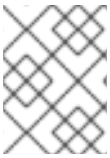
All users who want to use java-1.7.0-openjdk should install these newly released packages, which add this enhancement.

4.8. RHEA-2012:0838 – NEW PACKAGE: JAVA-1.7.0-ORACLE

New java-1.7.0-oracle package is now available for Red Hat Enterprise Linux 6.

The java-1.7.0-oracle package provides the Oracle Java 7 Runtime Environment and the Oracle Java 7 Software Development Kit.

This update adds the java-1.7.0-oracle packages to Red Hat Enterprise Linux 6. (BZ#720928)



NOTE

Before applying this update, make sure that any previous Oracle Java packages have been removed.

All users who require java-1.7.0-oracle should install these new packages.

4.9. RHEA-2012:1038 – NEW PACKAGES: KMOD-BNX2X, KMOD-BNX2, KMOD-BNX2I, KMOD-BNX2FC

New kmod-bnx2x, kmod-bnx2, kmod-bnx2i, kmod-bnx2fc packages are now available for Red Hat Enterprise Linux 6.

The kmod-bnx2x packages provide temporary drivers for the following hardware beyond what was delivered in Red Hat Enterprise Linux 6.2:

Broadcom NetXtreme II BCM57xx Gigabit Ethernet

The kmod-bnxx packages provide temporary drivers for the following hardware beyond what was delivered in Red Hat Enterprise Linux 6.2:

Broadcom NetXtreme II BCM5771x/578xx 10/20-Gigabit Ethernet

The kmod-bnx2i packages provide temporary drivers for the following hardware beyond what was delivered in Red Hat Enterprise Linux 6.2:

Broadcom NetXtreme II BCM570x/5771x/578xx iSCSI

The kmod-bnx2fc packages provide temporary drivers for the following hardware beyond what was delivered in Red Hat Enterprise Linux 6.2:

Broadcom NetXtreme II BCM5771x/578xx FCoE

This enhancement update adds the kmod-bnx2x, kmod-bnx2, kmod-bnx2i and kmod-bnx2fc packages to Red Hat Enterprise Linux 6 as part of the Red Hat Enterprise Linux Driver Update Program (DUP). The packages introduced by the RHEA-2012:0503 advisory did not contain proper firmware for kmod-

bnx2. In addition, the driver included in the kmod-bnx2x packages could, under certain circumstances, work incorrectly with Fibre Channel over Ethernet (FCoE). This update addresses these problems. (BZ#818940, BZ#819566, BZ#819569, BZ#819567)

Users encountering the aforementioned problems and users requiring temporary driver support for the specific hardware noted above should install these packages. Unless a system includes the exact hardware supported by kmod-bnx2x, kmod-bnx2, kmod-bnx2i, or kmod-bnx2fc, these packages must not be installed.

4.10. RHEA-2012:1576 – NEW PACKAGES: KMOD-PCH_GBE

New kmod-pch_gbe packages are now available for Red Hat Enterprise Linux 6.

The kmod-pch_gbe packages provide kernel modules for controlling Ethernet adapter in Intel EG20T Platform Controller Hub and OKI Semiconductor ML7223 Input/Output Hub.

The kmod-pch_gbe packages provide temporary drivers for the following hardware beyond what was delivered in Red Hat Enterprise Linux 6.3:

* Intel EG20T PCH / OKI Semiconductor ML7223 IOH Gigabit Ethernet

This enhancement update adds the kmod-pch_gbe packages to Red Hat Enterprise Linux 6 as part of the Red Hat Enterprise Linux Driver Update Program (DUP). (BZ#878375)

Only users requiring temporary driver support for the specific hardware noted above should install these packages. Unless a system includes the exact hardware explicitly supported by kmod-pch_gbe packages, these packages must not be installed.

4.11. RHEA-2012:0825 – NEW PACKAGE: LEDMON

A new ledmon package is 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.

This enhancement update adds the ledmon package to Red Hat Enterprise Linux 6. (BZ#750379)

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

4.12. RHEA-2012:0812 – NEW PACKAGE: LIBQB

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

The libqb package provides 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.

This enhancement update adds the libqb package to Red Hat Enterprise Linux 6. This package is introduced as a dependency of the pacemaker package, and is considered a Technology Preview in Red Hat Enterprise Linux 6.3. (BZ#782240)

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

4.13. RHEA-2012:0798 – NEW PACKAGES: LIBREOFFICE

New libreoffice packages 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.

This enhancement update adds the libreoffice packages to Red Hat Enterprise Linux 6. (BZ#747431)

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

4.14. RHEA-2012:0868 – NEW PACKAGES: LIBWACOM

New libwacom packages 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, calibrations, and provide device-specific configurations.

This enhancement update adds the libwacom packages to Red Hat Enterprise Linux 6. (BZ#786100)

All users who require libwacom should install these new packages.

4.15. RHEA-2012:0890 – NEW PACKAGE: NUMAD

A new numad package is now available as a Technology Preview for Red Hat Enterprise Linux 6.

The numad package provides a daemon for NUMA (Non-Uniform Memory Architecture) systems, that monitors NUMA characteristics. As an alternative to manual static CPU pining and memory assignment, numad provides dynamic adjustment to minimize memory latency on an ongoing basis. The package also provides an interface that can be used to query the numad daemon for the best manual placement of an application.

This enhancement update adds the numad package to Red Hat Enterprise Linux 6 as a Technology preview. (BZ#758416, BZ#824067)

More information about Red Hat Technology Previews is available here:

https://access.redhat.com/support/offerings/techpreview/

All users who want to use the numad Technology Preview should install this newly-released package, which adds this enhancement.

4.16. RHEA-2012:0826 – NEW PACKAGE: PPC64-DIAG

A new ppc64-diag package is now available for Red Hat Enterprise Linux 6.

The ppc64-diag package provides platform diagnostics for Linux for 64-bit PowerPC architectures.

This enhancement update adds the ppc64-diag package to Red Hat Enterprise Linux 6. (BZ#632735)

All users who require ppc64-diag are advised to install this newly released package.

4.17. RHEA-2012:0806 – NEW PACKAGES: SCL-UTILS

New scl-utils packages 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.

This enhancement update adds the scl-utils packages to Red Hat Enterprise Linux 6. (BZ#713147)

All users who require scl-utils should install these new packages.

4.18. RHEA-2012:0823 – NEW PACKAGE: SUBSCRIPTION-MANAGER-MIGRATION-DATA

A new subscription-manager-migration-data package is now available for Red Hat Enterprise Linux 6.

The new Subscription Management tooling allows users to understand the specific products, which have been installed on their machines, and the specific subscriptions, which their machines consume.

This enhancement update adds the subscription-manager-migration-data package to Red Hat Enterprise Linux 6. The package allows for migrations from Red Hat Network Classic Hosted to hosted certificate-based subscription management. (BZ#773030)

All users who require subscription-manager-migration-data are advised to install this new package.

4.19. RHEA-2012:0853 – NEW PACKAGES: USBREDIR

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

The usbredir packages provide a protocol for redirection of USB traffic from a single USB device to a different virtual machine then the one to which the USB device is attached. The usbredir package contains a number of libraries to help implement support for usbredir.

This enhancement update adds the usbredir package to Red Hat Enterprise Linux 6. (BZ#758098)

Users who wish to use the new USB redirection for Spice are advised to install these new packages.

4.20. RHEA-2012:0965 – NEW PACKAGE: VIRT-P2V

A new virt-p2v package is now available for Red Hat Enterprise Linux 6.

Virt-P2V is a tool for conversion of a physical server to a virtual guest.

This enhancement update adds the virt-p2v package to Red Hat Enterprise Linux 6, which contains a bootable ISO image for Virt-P2V conversion. The ISO image is also available on Red Hat Network in the Downloads section of the following channels:

- RHEL AUS Server (v. 6.2 for 64-bit x86_64)
- RHEL EUS Server (v. 6.2.z for 64-bit x86_64)
- Red Hat Enterprise Linux Client (v. 6 for 64-bit x86_64)

- Red Hat Enterprise Linux Compute Node (v. 6 for x86_64)
- Red Hat Enterprise Linux Server (v. 6 for 64-bit x86_64)
- Red Hat Enterprise Linux Workstation (v. 6 for x86_64)

The bootable image is needed to use the tool as the disks must be unmounted and not in use at the time of conversion. It allows you to convert servers running Microsoft Windows or Red Hat Enterprise Linux into virtual guests on Red Hat Enterprise Virtualization or libvirt hosts. For further information, refer to the V2V Guide. (BZ#807445)

All users who require virt-p2v should install this new package.

4.21. RHEA-2012:0786 – NEW PACKAGES: KMOD-HPSA

New kmod-hpsa packages are now available for Red Hat Enterprise Linux 6.

The kmod-hpsa packages provide kernel modules for controlling HP Smart Array Controllers.

The kmod-hpsa packages provide temporary drivers for the following hardware beyond what was delivered in Red Hat Enterprise Linux 6.3:

- HP Smart Array Controllers

This enhancement update adds the kmod-hpsa packages to Red Hat Enterprise Linux 6 as part of the Red Hat Enterprise Linux Driver Update Program (DUP).

Only users requiring temporary driver support for the specific hardware noted above should install these packages. Unless a system includes the exact hardware explicitly supported by the kmod-hpsa packages, these packages must not be installed.

Note that before installation of Red Hat Enterprise Linux 6.3 with the hpsa Driver Update Disk (DUD), you are advised to use the "pci=nomsi" kernel parameter to work around the hpsa driver load/unload issue described in BZ#904945. Once the installation is complete, this kernel parameter is no longer needed.

CHAPTER 5. PACKAGE UPDATES

5.1. 389-DS-BASE

5.1.1. RHBA-2012:1067 – 389-ds-base bug fix update

Updated 389-ds-base packages that fix two 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#834096

Prior to this update, simultaneous updates that included deleting an attribute in an entry could cause the domain directory server to abort with a segmentation fault. This update checks whether a modified attribute entry has a NULL value. Now, the server handles simultaneous updates as expected.

BZ#836251

Prior to this update, the get_entry function did not accept a NULL pblock. As a consequence, the Account Usability feature did not return the correct information about user account expiration and locked status. This update modifies the underlying code so that the get_entry function now accepts a NULL pblock.

All users of 389-ds-base are advised to upgrade to these updated packages, which fix these bugs. Note: after completing this update, the 389 server service is restarted automatically.

5.1.2. RHSA-2012:0997 – Moderate: 389-ds-base security update

Updated 389-ds-base packages that fix two security issues 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 the descriptions below.

The 389 Directory Server is an LDAPv3 compliant server. The 389-ds-base packages include the Lightweight Directory Access Protocol (LDAP) server and command line utilities for server administration.

Security Fixes

CVE-2012-2678

A flaw was found in the way 389 Directory Server handled password changes. If an LDAP user has changed their password, and the directory server has not been restarted since that change, an attacker able to bind to the directory server could obtain the plain text version of that user's password via the "unhashed#user#password" attribute.

CVE-2012-2746

It was found that when the password for an LDAP user was changed, and audit logging was enabled

(it is disabled by default), the new password was written to the audit log in plain text form. This update introduces a new configuration parameter, "nsslapd-auditlog-logging-hide-unhashed-pw", which when set to "on" (the default option), prevents 389 Directory Server from writing plain text passwords to the audit log. This option can be configured in **/etc/dirsrv/slapd-*ID*/dse.ldif**.

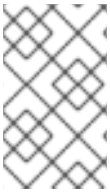
All users of 389-ds-base are advised to upgrade to these updated packages, which resolve these issues. After installing this update, the 389 server service will be restarted automatically.

5.1.3. RHSA-2012:0813 — Low: 389-ds-base security, bug fix and enhancement update

Updated 389-ds-base 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 score, which gives a detailed severity rating, is available from the CVE link associated with the description below.

The 389-ds-base packages provide 389 Directory Server, which is an LDAPv3 (Lightweight Directory Access Protocol version 3) compliant server, and command-line utilities for server administration.



NOTE

The 389-ds-base package has been upgraded to upstream version 389-ds-base-1.2.10, which provides a number of bug fixes and enhancements over the previous version. (BZ#766989)

Security Fix

CVE-2012-0833

A flaw was found in the way the 389 Directory Server daemon (ns-slapd) handled access control instructions (ACIs) using certificate groups. If an LDAP user that had a certificate group defined attempted to bind to the directory server, it would cause ns-slapd to enter an infinite loop and consume an excessive amount of CPU time.

Red Hat would like to thank Graham Leggett for reporting this issue.

Bug Fixes

BZ#743979

Previously, 389 Directory Server used the Netscape Portable Runtime (NSPR) implementation of the read/write locking mechanism. Consequently, the server sometimes stopped responding to requests under heavy loads. This update replaces the original locking mechanism with the POSIX (Portable Operating System Interface) read/write locking mechanism. The server is now always responsive under heavy loads.

BZ#745201

Previously, Distinguished Names (DNs) were not included in access log records of LDAP compare operations. Consequently, this information was missing in the access logs. This update modifies the underlying source code so that DNs are logged and can be found in the access logs.

BZ#752577

Previously, when 389 Directory Server was under heavy load and operating in a congested network, problems with client connections sometimes occurred. When there was a connection problem while the server was sending Simple Paged Result (SPR) search results to the client, the LDAP server called a cleanup routine incorrectly. Consequently, a memory leak occurred and the server terminated unexpectedly. This update fixes the underlying source code to ensure that cleanup tasks are run correctly and no memory leaks occur. As a result, the server does not terminate or become unresponsive under heavy loads while servicing SPR requests.

BZ#757897

Previously, certain operations with the Change Sequence Number (CSN) were not performed efficiently by the server. Consequently, the **ns-slapd** daemon consumed up to 100% of CPU time when performing a large number of CSN operations during content replication. With this update, the underlying source code has been modified to perform the CSN operations efficiently. As a result, large numbers of CSN operations can be performed during content replications without any performance issues.

BZ#757898

Previously, allocated memory was not correctly released in the underlying code for the SASL GSSAPI authentication method when checking the Simple Authentication and Security Layer (SASL) identity mappings. This problem could cause memory leaks when processing SASL bind requests, which eventually caused the LDAP server to terminate unexpectedly with a segmentation fault. This update adds function calls that are needed to free allocated memory correctly. Memory leaks no longer occur and the LDAP server no longer crashes in this scenario.

BZ#759301

Previously, 389 Directory Server did not handle the Entry USN (Update Sequence Number) index correctly. Consequently, the index sometimes became out of sync with the main database and search operations on USN entries returned incorrect results. This update modifies the underlying source code of the Entry USN plug-in. As a result, the Entry USN index is now handled by the server correctly.

BZ#772777

Previously, search filter attributes were normalized and substring regular expressions were compiled repeatedly for every entry in the search result set. Consequently, using search filters with many attributes and substring subfilters resulted in poor search performance. This update ensures that search filters are pre-compiled and pre-normalized before being applied. These changes result in better search performance when applying search filters with many attributes and substring subfilters.

BZ#772778

Previously, the number of ACIs (Access Control Information records) to be cached was limited to 200. Consequently, evaluating a Directory Server entry against more than 200 ACIs failed with the following error message:

```
acl_TestRights - cache overflown
```

This update increases the default ACI cache limit to 2000 and allows it to be configurable by means of the new parameter ***nsslapd-aclpb-max-selected-acls*** in the configuration file entry "cn=ACL Plugin,cn=plugins,cn=config". As a result, the aforementioned error message is not displayed unless the new limit is exceeded, and it is now possible to change the limit when needed.

BZ#772779

Previously, the restore command contained a code path leading to an infinite loop. Consequently, 389 Directory Server sometimes became unresponsive when performing a restore from a database backup. This update removes the infinite loop code path from the underlying source code. As a result, the server does not stop responding when performing a database restore.

BZ#781485

Previously, performing the **ldapmodify** operation to modify RUV (Replica Update Vector) entries was allowed. Consequently, 389 Directory Server became unresponsive when performing such operations. This update disallows direct modification of RUV entries. As a result, the server does not stop responding when performing such operations, and returns an error message advising usage of the **CLEANRUV** operation instead.

BZ#781495

Previously, to identify restart events of 389 Directory Server, the **logconv.pl** script searched server logs for the "conn=0 fd=" string. Consequently, the script reported a wrong number of server restarts. This update modifies the script to search for the "conn=1 fd=" string instead. As a result, the correct number of server restarts is now returned.

BZ#781500

When reloading a database from an LDIF (LDAP Data Interchange Format) file that contained an RUV element with an obsolete or decommissioned replication master, the changelog was invalidated. As a consequence, 389 Directory Server emitted error messages and required re-initialization. This update ensures that the user is properly informed about obsolete or decommissioned replication masters, and that such masters are deleted from the RUV entries. Database is now reloaded as expected in this scenario.

BZ#781516

Previously, when a non-leaf node became a tombstone entry, its child entries lost the parent-child relationships. Consequently, non-leaf tombstone entries could have been reaped prior to their child tombstone entries. This update fixes the underlying source code so that parent-child relationships are maintained even when a non-leaf entry is deleted. As a result, tombstones are now reaped correctly in the bottom-up order.

BZ#781529

Previously, no validation of managed entry attributes against the managed entry template was performed before updating 389 Directory Server's managed entries. Consequently, managed entries could have been updated after updating an original entry attribute that was not contained in the managed entry template. This update adds a check that compares modified attributes with managed entry template attributes. As a result, the managed entries are not updated unless the modified attributes of the original entry are contained in the managed entry template.

BZ#781533

Previously, 389 Directory Server did not shut down before all running tasks had been completed. Consequently, it sometimes took a long time for the Directory Server to shut down when a long-running task was being carried out. This update enhances the underlying source code with a check for server shutdown requests during performance of long-running tasks. As a result, the server shuts down in a standard amount of time even when a long-running task is being processed.

BZ#781537

Previously, 389 Directory Server expected the value of the ***authzid*** attribute to be fully BER (Basic Encoding Rules) encoded. Consequently, the following error was returned when performing the **ldapsearch** command with proxy authorization:

```
unable to parse proxied authorization control (2 (protocol error))
```

This update modifies the underlying source code so that full BER encoding of the provided authzid value is not required. As a consequence, no error is returned in the scenario described above.

BZ#781538

Previously, the buffer for matching rule OIDs (Object Identifiers) had a fixed size of 1024 characters. Consequently, matching rule OIDs got truncated when their total length exceeded 1024 characters. This update modifies the underlying source code to use a dynamically allocated buffer instead of the one with a fixed size. As a result, any number of matching rule OIDs can be handled without being truncated.

BZ#781539

Previously, executing the **ldapsearch** command on the "cn=config" object returned all attributes of the object, including attributes with empty values. This update ensures that attributes with empty values are not saved into "cn=config", and enhances the **ldapsearch** command with a check for empty attributes. As a result, only attributes that have a value are returned in the aforementioned scenario.

BZ#781541

Previously, log records of operations performed using a proxy user contained the main user as the one who performed the operation. This update ensures that the proxy user is logged in log records of the search, add, mod, del, and modrdn operations.

BZ#784343

Previously, the database upgrade scripts checked if the server was offline by checking for the presence of **.pid** files. In some cases, however, the files remain present even if the associated processes have already been terminated. Consequently, the upgrade scripts sometimes assumed that the Directory Server was online and did not proceed with the database upgrade even if the server was actually offline. This update adds an explicit test to check if the processes referenced in the **.pid** files are really running. As a result, the upgrade scripts now work as expected.

BZ#784344

Previously, the **repl-monitor** command used only the subdomain part of hostnames for host identification. Consequently, hostnames with the identical subdomain part (for example: "ldap.domain1", "ldap.domain2") were identified as a single host, and inaccurate output was produced. This update ensures that the entire hostname is used for host identification. As a result, all hostnames are identified as separate and output of the **repl-monitor** command is accurate.

BZ#788140

Previously, the server used unnormalized DN strings to perform internal search and modify operations while the code for modify operations expected normalized DN strings. Consequently, error messages like the following one were logged when performing replication with domain names specified in unnormalized format:

```
NSMMReplicationPlugin - repl_set_mtn_referrals: could not set referrals for
replica dc=example,dc=com: 32
```

This update ensures that DN strings are normalized before being used in modify operations. As a result, replication does not produce the error messages in the aforementioned scenario.

BZ#788722

Previously, the **389-ds-base/ldap/servers/snmp/** directory contained **.mib** files without copyright headers. Consequently, the files could not be included in certain Linux distributions due to copyright reasons. This update merges information from all such files into the **redhat-directory.mib** file, which contains the required copyright information, and ensures that it is the only file in the directory. As a result, no copyright issues block 389 Directory Server from being included in any Linux distribution.

BZ#788724

Previously, the underlying source code for extensible search filters used **strcmp** routines for value comparison. Consequently, using extensible search filters with binary data returned incorrect results. This update modifies the underlying source code to use binary-aware functions. As a result, extensible search filters work with binary data correctly.

BZ#788725

Previously, value normalization of the search filter did not respect the used filter type and matching rules. Consequently, when using different values than the default comparison type for the searched attribute syntax, search attempts returned incorrect results. This update modifies the underlying source code to use normalization sensitive to matching rules on filter attributes and values. As a result, search results in accordance with the matching rules are returned.

BZ#788729

Previously on the Directory Server, tombstones of child entries in a database were handled incorrectly. Therefore, if the database contained deleted entries that were converted to tombstones, an attempt to reindex the **entryrdn** index failed with the following error message:

```
_entryrdn_insert_key: Getting "nsuniqueid=ca681083-69f011e0-8115a0d5-
f42e0a24,ou=People,dc=example,dc=com" failed
```

With this update, 389 Directory Server handles tombstones of child entries correctly, and the **entryrdn** index can now be reindexed successfully with no errors.

BZ#788731

Previously, RUV tombstone entries were indexed incorrectly by the **entryrdn** index. Consequently, attempts to search for such entries were not successful. This update ensures correct indexing of RUV tombstone entries in the **entryrdn** index and search attempts for such entries are now successful.

BZ#788741

Previously, the DNA (Distributed Numeric Assignment) plug-in used too short timeout for requests to replicate a range of UIDs. Consequently, using replication with DNA to add users sometimes failed on networks with high latency, returning the following error message:

```
Operations error: Allocation of a new value for range cn=posix ids,cn=distributed
numeric assignment plugin,cn=plugins,cn=config failed
```

With this update, the default timeout for such replication requests has been set to 10 minutes. As a result, no errors are returned when using replication with DNA to add users, and the operation succeeds.

BZ#788745

Previously, change sequence numbers (CSNs) in RUV were not refreshed when a replication role was changed. Consequently, data on the server became inconsistent. This update ensures that CSNs are

refreshed when a replication role is changed. As a result, data inconsistency is no longer observed in the previously mentioned cases.

BZ#788749

Previously, errors in schema files were not reported clearly in log files. Consequently, the messages could be incorrectly interpreted as reporting an error in the **dse.ldif** file. This update modifies the error messages so that they include the name of and path to the file where the error was found.

BZ#788750

Previously, the server used an outdated version of the nisDomain schema after an upgrade. Consequently, restarting 389 Directory Server after an upgrade produced the following error message:

```
attr_syntax_create - Error: the EQUALITY matching rule [caseIgnoreMatch] is not
compatible with the syntax [1.3.6.1.4.1.1466.115.121.1.26] for the attribute [nisDomain]
```

This update ensures that the server uses the latest version of the nisDomain schema. As a result, restarting the server after an upgrade does not show any errors.

BZ#788751

389 Directory Server previously did not properly release allocated memory after finishing normalization operations. This caused memory leaks to occur during server's runtime. This update fixes the underlying code to release allocated memory properly so that memory leaks no longer occur under these circumstances

BZ#788753

Previously, the "connection" attribute was not included in the cn=monitor schema, which caused the access control information (ACI) handling code to ignore the ACI. Consequently, requesting the ***connection*** attribute when performing anonymous search on cn=monitor returned the ***connection*** attribute, even though it was denied by the default ACI. This update ensures that the ACI is processed even if the attribute is not in the schema. As a result, the ***connection*** attribute is not displayed if the ACI denies it.

BZ#788754

Previously, several memory leak errors sometimes occurred during the server's runtime. This update fixes all the memory leak errors so that none of them occur anymore.

BZ#788755

Previously, IPv4-mapped IPv6 addresses were treated as independent addresses by 389 Directory Server. Consequently, errors were reported during server startup when such addresses conflicted with standard IPv4 addresses. This update ensures that the IPv4 part of every IPv4-mapped IPv6 address is compared with existing IPv4 addresses. As a result, the server starts with no errors even when IPv4-mapped IPv6 addresses conflict with standard IPv4 addresses.

BZ#788756

Previously, the 389-ds-base man pages contained several typos and factual errors. This update corrects the man pages so that they contain correct information and no typos.

BZ#790491

Previously, a NULL pointer dereference sometimes occurred when initializing a Directory Server replica. Consequently, the server terminated unexpectedly with a segmentation fault. This update

enhances the underlying source code for replica initialization with a check for the NULL value. As a result, replica initialization always finishes successfully.

BZ#796770

Previously, a double free error sometimes occurred during operations with orphaned tombstone entries. Consequently, when an orphaned tombstone entry was passed to the **tombstone_to_glue** function, the Directory Server terminated unexpectedly. This update fixes the logic for getting ancestor tombstone entries and eliminates the chance to convert a tombstone entry into an orphaned entry. As a result, unexpected server termination no longer occurs in the aforementioned scenario.

BZ#800215

Previously, an internal loop was incorrectly handled in code of the **ldapcompare** command. Consequently, performing concurrent comparison operations on virtual attributes caused the Directory Server to become unresponsive. This update fixes the internal loop issue. As a result, the server performs concurrent comparison operations without any issues.

BZ#803930

Previously, when upgrading 389 Directory Server, server startup had been initiated before the actual upgrade procedure finished. Consequently, the startup failed with the following error message:

```
ldif2dbm - _get_and_add_parent_rdns: Failed to convert DN cn=TESTRELM.COM to RDN
```

This update ensures that the server does not start before the upgrade procedure finishes. As a result, the server boots up successfully after the upgrade.

BZ#811291

Previously, the code of the range read operation did not correctly handle situations when an entry was deleted while a ranged search operation was being performed. Consequently, performing delete and ranged search operations concurrently under heavy loads caused the Directory Server to terminate unexpectedly. This update fixes the underlying source code to handle such situations correctly. As a result, the server does not terminate before performing delete and ranged search operations concurrently under heavy loads.

BZ#813964

When performing delete and search operations against 389 Directory Server under high load, the DB_MULTIPLE_NEXT pointer to the stack buffer could have been set to an invalid value. As a consequence, pointer's dereference lead to an attempt to access memory that was not allocated for the stack buffer. This caused the server to terminate unexpectedly with a segmentation fault. With this update, the DB_MULTIPLE_NEXT pointer is now properly tested. If the pointer's value is invalid, the page or value is considered deleted and the stack buffer is reloaded. As a result, the segmentation fault no longer occurs in this scenario.

BZ#815991

The **ldap_initialize()** function is not thread-safe. Consequently, 389 Directory Server terminated unexpectedly during startup when using replication with many replication agreements. This update ensures that calls of the **ldap_initialize()** function are protected by a mutual exclusion. As a result, when using replication with many replication agreements, the server starts up correctly.

BZ#819643

Due to an error in the underlying source code, an attempt to rename an RDN (Relative Distinguished

Name) string failed if the new string sequence was the same except of using the different lower/upper case of some letters. This update fixes the code so that it is possible to rename RDNs to the same string sequence with case difference.

BZ#821542

Previously, the letter case information was ignored when renaming DN strings. Consequently, if the new string sequence differed only in the case of some letters, a DN string was only converted to lowercase and the case information lost. This update modifies the underlying code so that it is now possible to rename RDNs to the same string sequence with case difference.

BZ#822700

Previously, the code for ACI handling did not reject incorrectly specified DNs. Consequently, incorrectly specified DNs in an ACI caused 389 Directory Server to terminate unexpectedly during startup or after an online import. This update ensures that the underlying source code for ACI handling rejects incorrectly specified DNs. As a result, the server does not terminate in this scenario.

BZ#824014

Previously, the code handling the "**entryusn**" attribute modified cache entries directly. Consequently under heavy loads, the server terminated unexpectedly when performing delete and search operations using the "**entryusn**" and "**memberof**" attributes with referential integrity enabled. This update ensures that the entries are never modified in the cache directly. As a result, the server performs searches in the previously described conditions without terminating unexpectedly.

Enhancements

BZ#683241

Previously, post-operation plug-ins were executed after initial operation results had been returned to the LDAP client. Consequently, some results of the initial operation might not have been immediately available. This update introduces the "betxnpreoperation" and "betxnpostoperation" plug-in types. Plug-ins of these types run inside the regular transaction of initial operations. As a result, when these plug-in types are used, operations triggered by the initial operation complete before completion of the initial operation.

BZ#766322

Previously, there was no easy way to determine what default search base an LDAP client should use. Consequently, LDAP clients with no search base configured attempted to search against 389 Directory Server. This update adds a new attribute, defaultNamingContext, to the root DSE (Directory Server Entry). As a result, clients can query the root DSE for the value of the defaultNamingContext attribute and use the returned value as a search base.

BZ#768086

This update introduces the nsslapd-minssf-exclude-rootdse configuration attribute, with possible values "on" and "off". If its value is "off", which is the default, the server allows clients to access the root DSE even if the Security Strenght Factor (SSF) value is less than the nsslapd-minssf attribute value. As a result, it is possible to allow access to the root DSE without using SSL/TLS even if the rest of the server requires SSL/TLS.

BZ#768091

Previously, the delete operation was not allowed for Managed Entry Config entries. Consequently, attempts to delete such entries were rejected with the following error message:

```
ldap_delete: Server is unwilling to perform (53)
additional info: Not a valid operation.
```

This update modifies the underlying source code so that deletion of Managed Entry Config entries is allowed and can be performed successfully.

BZ#781501

Previously, extended user account information was not available to LDAP clients from 389 Directory Server. This update adds support for Account Usable Request Control, which enables LDAP clients to get the extended user account information.

BZ#788760

Previously, the **logconv.pl** script was only able to produce a summary of operations for a file or for a requested period. This update introduces the **-m** option for generation of per-second statistics, and the **-M** option for generation of per-minute statistics. The statistics are generated in CSV format suitable for further post-processing.

BZ#790433

Previously, all newly created entries had to be added to groups manually. This update adds a new plug-in which ensures automatic adding of each new entry to a group if it matches certain criteria.

Users of 389-ds-base should upgrade to these updated packages, which resolve these issues and add these enhancements.

5.2. ABRT AND LIBREPORT

5.2.1. RHSA-2013:0215 – Important: abrt and libreport security update

Updated abrt and libreport packages that fix two security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

ABRT (Automatic Bug Reporting Tool) is a tool to help users to detect defects in applications and to create a bug report with all the information needed by a maintainer to fix it. It uses a plug-in system to extend its functionality. libreport provides an API for reporting different problems in applications to different bug targets, such as Bugzilla, FTP, and Trac.

Security Fixes

CVE-2012-5659

It was found that the /usr/libexec/abrt-action-install-debuginfo-to-abrt-cache tool did not sufficiently sanitize its environment variables. This could lead to Python modules being loaded and run from non-standard directories (such as /tmp/). A local attacker could use this flaw to escalate their privileges to that of the abrt user.

CVE-2012-5660

A race condition was found in the way ABRT handled the directories used to store information about crashes. A local attacker with the privileges of the abrt user could use this flaw to perform a symbolic link attack, possibly allowing them to escalate their privileges to root.

Red Hat would like to thank Martin Carpenter of Citco for reporting the CVE-2012-5660 issue. CVE-2012-5659 was discovered by Miloslav Trmač of Red Hat.

All users of abrt and libreport are advised to upgrade to these updated packages, which correct these issues.

5.3. ABRT, LIBREPORT, BTPARSER, AND PYTHON-MEH

5.3.1. RHSA-2012:0841 – Low: abrt, libreport, btparser and python-meh security and bug fix update

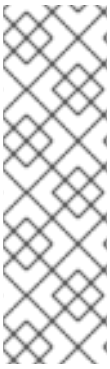
Updated abrt, libreport, btparser, and python-meh packages that fix two security issues and several bugs are now available for Red Hat Enterprise Linux 6.

The Red Hat Security Response Team has rated this update as having 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.

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. It uses a plug-in system to extend its functionality. libreport provides 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.

The python-meh package provides a python library for handling exceptions.



NOTE

The abrt package has been upgraded to upstream version 2.0.8-1, which provides a number of bug fixes over the previous version. (BZ#759375)

The libreport package has been upgraded to upstream version 2.0.9-1, which provides a number of bug fixes over the previous version. (BZ#759377)

The btparser package has been upgraded to upstream version 0.16-1, which provides a number of bug fixes over the previous version. (BZ#768377)

Security Fixes

CVE-2012-1106

If the C handler plug-in in ABRT was enabled (the abrt-addon-ccpp package installed and the abrt-ccpp service running), and the sysctl fs.suid_dumpable option was set to "2" (it is "0" by default), core dumps of set user ID (setuid) programs were created with insecure group ID permissions. This could allow local, unprivileged users to obtain sensitive information from the core dump files of setuid processes they would otherwise not be able to access.

CVE-2011-4088

ABRT did not allow users to easily search the collected crash information for sensitive data prior to submitting it. This could lead to users unintentionally exposing sensitive information via the submitted crash reports. This update adds functionality to search across all the collected data. Note that this fix does not apply to the default configuration, where reports are sent to Red Hat Customer Support. It only takes effect for users sending information to Red Hat Bugzilla.

Red Hat would like to thank Jan Iven for reporting CVE-2011-4088.

Bug Fixes

BZ#809587, BZ#745976

When the ABRT GUI was used to report a bug using the menu button **Report problem with ABRT**, an empty bug was created. This update removes this button as it was only used for testing purposes.

BZ#800828

When a new dump directory was saved to **/var/spool/abrt-upload/** via the **reporter-upload** utility, the ABRT daemon copied the dump directory to **/var/spool/abrt/** and incremented the crash count which was already incremented before. Due to the crash count being incremented twice, the dump directory was marked as a duplicate of itself and removed. With this update, the crash count is no longer incremented for remotely uploaded dump directories, thus fixing the issue.

BZ#747624

The **/usr/bin/abrt-cli** utility was missing a man page. This update adds the **abrt-cli(1)** man page.

BZ#796216

Analyzing lines of a kernel oops caused the **line** variable to be freed twice. This update fixes this bug, and kernel oopses are now properly analyzed.

BZ#770357

Prior to this update, ABRT email notification via the **mailx** plug-in did not function properly due to a missing default configuration file for the **mailx** plug-in. This update adds a default configuration file for the **mailx** plug-in: **/etc/libreport/plugins/mailx.conf**.

BZ#799352

Starting the ABRT daemon resulted in an error if **dbus** was not installed on the system. This update removes the **dbus** dependency and the ABRT daemon can now be started even if **dbus** is not installed on the system.

BZ#727494

The previous version of ABRT silently allowed users to report the same problem to Bugzilla multiple times. This behavior is now changed and users are warned if the report was already submitted. The max allowed size of email attachments and local logs was increased to 1 MB. This fixes the problem where longer reports were being lost when sent via email or stored locally using the **logger** plug-in.

BZ#746727

This update fixes a bug which caused the **/tmp/anaconda-tb-*** files to be sometimes recognized as a binary file and sometimes as a text file.

BZ#771597

ABRT 2.x has added various new daemons. However, not all of the added daemons were properly enabled during the transition from ABRT 1.x. With this update, all daemons are correctly started and updating from ABRT 1.x to ABRT 2.x works as expected.

BZ#751068

The abrt-cli package previously depended on the abrt-addon-python package. This prevented users from removing the abrt-addon-python package via Yum as the abrt-cli would be removed as well. With this update, a new "virtual" abrt-tui package has been added that pulls all the required packages in order to use ABRT on the command line, thus, resolving the aforementioned issue.

BZ#749100

Previously, some strings in the ABRT tools were not marked as translatable. This update fixes this issue.

BZ#773242

When ABRT attempted to move data, a misleading message was returned to the user informing that a copy of the dump was created. This update improves this message so that it is clear that ABRT does not copy data but moves it.

BZ#811147

When a backtrace contains a frame with text consisting of function arguments that was too long, the backtrace printer in GDB truncates the arguments. The backtrace parser could not handle the truncated arguments and did not format them properly. With this update, the backtrace parser detects the truncated strings, indicating the function arguments were truncated. The parser state then adapts to this situation and correctly parses the backtrace.

BZ#823411

A change in the Bugzilla API prevented the ABRT **bugzilla** plug-in from working correctly. This update resolves this issue by modifying the source code to work with the new Bugzilla API.

BZ#758366

This update fixes a typographical error in the commentary of various ABRT configuration files.

BZ#625485

The previous version of ABRT generated an invalid XML log file. This update fixes this and every non-ASCII character is now escaped.

BZ#788577

Unlike ABRT, **python-meh** was not including a list of environment variables in its problem reports. A list of environment variables is useful information for assignees of the created bug. With this update, code producing a list of environment variables and passing it to **libreport** was added to **python-meh**, and problem reports generated by **python-meh** now include lists of environment variables.

All users of abrt, libreport, btparser, and python-meh are advised to upgrade to these updated packages, which contain backported patches to correct these issues.

5.4. ACROREAD

5.4.1. RHSA-2013:0150 – Critical: acroread security update

Updated acroread packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

Adobe Reader allows users to view and print documents in Portable Document Format (PDF).

Security Fix

CVE-2012-1530, CVE-2013-0601, CVE-2013-0602, CVE-2013-0603, CVE-2013-0604, CVE-2013-0605, CVE-2013-0606, CVE-2013-0607, CVE-2013-0608, CVE-2013-0609, CVE-2013-0610, CVE-2013-0611, CVE-2013-0612, CVE-2013-0613, CVE-2013-0614, CVE-2013-0615, CVE-2013-0616, CVE-2013-0617, CVE-2013-0618, CVE-2013-0619, CVE-2013-0620, CVE-2013-0621, CVE-2013-0623, CVE-2013-0626

This update fixes several security flaws in Adobe Reader. These flaws are detailed in the Adobe Security bulletin APSB13-02. A specially-crafted PDF file could cause Adobe Reader to crash or, potentially, execute arbitrary code as the user running Adobe Reader when opened.

All Adobe Reader users should install these updated packages. They contain Adobe Reader version 9.5.3, which is not vulnerable to these issues. All running instances of Adobe Reader must be restarted for the update to take effect.

5.5. ALSA-UTILS

5.5.1. RHBA-2012:0917 – alsa-utils bug fix and enhancement update

Updated alsa-utils packages that fix one bug and add one enhancement are now available for Red Hat Enterprise Linux 6.

The alsa-utils packages provide command-line utilities for the Advanced Linux Sound Architecture (ALSA).

Bug Fix

BZ#674199

Prior to this update, the alsactl tool tried to initialize all sound cards if the /etc/asound.state file was not present. As a consequence, SElinux could deny access to non-existent devices. This update modifies the underlying code so that alsactl is called only once from udev.

Enhancement

BZ#650113

With this update, the alsa-delay and alsaloop utilities have been added to alsa-utils to manage the system audio delay.

All users of alsa-utils are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

5.6. ANACONDA

5.6.1. RHBA-2012:0782 – 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.

The anaconda package contains portions of the Anaconda installation program that can be run by the user for reconfiguration and advanced installation options.

Bug Fixes

BZ#690058

Prior to this update, the **noprobe** argument in a **kickstart** file was not passed to the last known codepath. Consequently, the **noprobe** request was not properly honored by **Anaconda**. This update improves the code so that the argument is passed to the last known codepath. As a result, device drivers are loaded according to the **device** command in the **kickstart** file.

BZ#691794

Previously, an improper device file that provided access to an array as a whole was used to initialize the boot loader in a *Device Mapper Multipath* (DM-Multipath) environment. Consequently, the system was not bootable. **Anaconda** has been modified to enumerate all drives in an array and initialize the boot loader on each of them. As a result, the system now boots as expected.

BZ#723404

When performing a minimal installation from media without the use of a network, network devices did not have a working default network configuration. Consequently, bringing a network device up after reboot using the **ifup** command failed. This update sets the value of ***BOOTPROTO*** to **dhcp** in default network device configuration files. As a result, network devices can be activated successfully using the **ifup** command after reboot in the scenario described.

BZ#727136

When **Anaconda** places a *PowerPC Reference Platform* (PReP) boot partition on a different drive to the root partition, the system cannot boot. This update forces the PReP boot partition to be on the same drive as the root partition. As a result, the system boots as expected.

BZ#734128

Due to a regression, when installing on systems with pre-existing mirrored *Logical Volumes* (LV), the installer failed to properly detect the **Logical Volume Management** configuration containing mirrored logical volumes. Consequently, a mirrored logical volume created before installation was not shown and could not be used in **kickstart**. The code to handle mirrored logical volumes has been updated to make use of the **udev** information that changed due to a previous bug fix. As a result, mirrored logical volumes are correctly detected by the installer.

BZ#736457

On IBM System z architectures, z/VM guests with only one CPU allocated failed to read the Conversational Monitor System (CMS) configuration file used by the installation environment. Consequently, users of z/VM guests with a single CPU had to either pass all installation environment configuration values on the kernel boot line or supply the information at the interactive prompts as the installation environment booted up. This update improves the code to detect the number of guests after mounting the /**proc** file. As a result, guests with one CPU can bring the boot device online so the CMS configuration file can be read and automated installations proceed as expected.

BZ#738577

The ***repo*** commands in **kickstart** generated by **Anaconda** contained base installation repository information but they should contain only additional repositories added either by the ***repo*** kickstart command or in the graphical user interface (GUI). Consequently, in media installations, the ***repo*** command generated for installation caused a failure when the kickstart file was used. With this update, Anaconda now generates ***repo*** commands only for additional repositories. As a result, kickstart will not fail for media installations.

BZ#740870

Manual installation on to BIOS RAID devices of level 0 or level 1 produced an *Intel Media Storage Manager* (IMSM) metadata read error in the installer. Consequently, users were not able to install to such devices. With this update, **Anaconda** properly detects BIOS RAID level 0 and level 1 IMSM metadata. As a result, users are able to install to these devices.

BZ#746495

The LiveCD environment was missing a legacy symlink to the **devkit-disks** utility. Consequently, the call that modified automounter behavior was never properly executed. The code has been updated to call the proper non-legacy binary. As a result, **USB** devices used during installation are no longer automounted.

BZ#747219

The console **tty1** was put under control of **Anaconda**, but was not returned when Anaconda exited. Consequently, **init** did not have permission to modify tty1's settings to enable **Ctrl**+**C** functionality when Anaconda exited, which resulted in **Ctrl**+**C** not working when the installer prompted the user to press the **Ctrl**+**C** or **Ctrl**+**Alt**+**Delete** key combination after Anaconda terminated unexpectedly. A code returning tty1 control back to init was added to Anaconda. As a result, **Ctrl**+**C** now works as expected if the user is prompted to press it when Anaconda crashes.

BZ#750126

The Bash version used in the **buildinstall** script had a bug that influenced parsing of the **=~** operator. This operator is used to check for the architecture when including files. Consequently, some binaries which provide the **grub** command were present on x86_64 versions of the installer, but were missing from i686 media. The Bash code has been modified to prevent this bug. As a result, the binaries are now also present on i686 media and users can now use the grub command from installation media as expected.

BZ#750417

Due to bad ordering in the unmounting sequence, the dynamic linker failed to link libraries, which caused the **mdadm** utility not to work and exit with the status code of **127**. This update fixes the ordering in the unmounting sequence and as a result, the dynamic linker and mdadm now work correctly.

BZ#750710

There was no check to see if the file descriptors passed as **stdout** and **stderr** were distinct. Consequently, if the stdout and stderr descriptors were the same, using them both for writing resulted in overwriting and the log file not containing all of the lines expected. With this update, if the stdout and stderr descriptors are the same then only one of them is used for both stdin and stderr. As a result, the log file contains all lines from both stdout and stderr.

BZ#753108

When installing on a system with more than one disk with a *PowerPC Reference Platform* (PReP) partition present, the PReP partitions that should be left untouched were updated. This update corrects the problem so that PReP partitions other than the one used during installation are left

untouched. As a result, old PReP partitions do not get updated.

BZ#754031

The kernel command line **/proc/cmdline** ends with **\n** but the installer only checked for **\0**. Consequently, the **devel** argument was not detected when it was the last argument on the command line and the installation failed. This update improves the code to also check for **\n**. As a result, the **devel** argument is correctly parsed and installation proceeds as expected.

BZ#756608

Network installations on IBM System z check the **nameserver** address provided using the ping command. Environments restricting **ICMP ECHO** packets will cause this test to fail, halting the installation and asking the user whether or not the provided nameserver address is valid. Consequently, automated installations using kickstart will stop if this test fails. With this update, in the event that the ping test fails, the **nslookup** command is used to validate the provided nameserver address. If the **nslookup** test succeeds then **kickstart** will continue with the installation. As a result, automated network installations on IBM System z in non-interactive mode will complete as expected in the scenario described.

BZ#760250

When configuring a system with multiple active network interfaces and the ***ksdevice*** = **link** command was present, the **link** specification was not used consistently for device activation and device configuration. Consequently, other network devices having link status were sometimes misconfigured using the settings targeted to the device activated by the installer. With this update, the code has been improved and now refers to the same device with **link** specification both in case of device activation and device configuration. As a result, when multiple devices with link status are present during installation, ***ksdevice*** = **link** specification of the device to be activated and used by the installer does not cause misconfiguration of another device having link status.

BZ#766902

When configuring the network using the **Anaconda** GUI hostname screen, the keyboard shortcut for the **Configure Network** button was missing. This update adds the **C** keyboard shortcut. Network configuration can now be invoked using the **Alt**+**C** keyboard shortcut.

BZ#767727

The Ext2FS class in **Anaconda** has a maximum file size attribute correctly set to **8 TB**, but Ext3FS and Ext4FS inherited this value without overriding it. Consequently, when attempting to create an ext3 or ext4 file system of a size greater than **8Tb** the installer would not allow it. With this update, the installer's upper bound for new ext3 and ext4 filesystem size has been adjusted from **8Tb** to **16TB**. As a result, the installer now allows creation of ext3 and ext4 filesystems up to **16TB**.

BZ#769145

The **Anaconda** dhcptimeout boot option was not working. **NetworkManager** used a **DHCP** transaction timeout of 45 seconds without the possibility of configuring a different value. Consequently, in certain cases NetworkManager failed to obtain a network address. **NetworkManager** has been extended to read the timeout parameter from a DHCP configuration file and use that instead of the default value. Anaconda has been updated to write out the dhcptimeout value to the interface configuration file used for installation. As a result, the boot option **dhcptimeout** works and NetworkManager now waits to obtain an address for the duration of the DHCP transaction period as specified in the DHCP client configuration file.

BZ#783245

Prior to this update, **USB3** modules were not in the **Anaconda** install image. Consequently, USB3 devices were not detected by Anaconda during installation. This update adds the USB3 modules to the install image and USB3 devices are now detected during installation.

BZ#783841

When the **kickstart clearpart** command or the installer's automatic partitioning options to clear old data from the system's disks were used with complex storage devices such as logical volumes and software RAID, **LVM** tools caused the installation process to become unresponsive due to a deadlock. Consequently, the installer failed when trying to remove old metadata from complex storage devices. This update changes the LVM commands in the **udev** rules packaged with the installer to use a less restrictive method of locking and the installer was changed to explicitly remove partitions from a disk instead of simply creating a new partition table on top of the old contents when it initializes a disk. As a result, LVM no longer hangs in the scenario described.

BZ#785400

The **/usr/lib/anaconda/textw/netconfig_text.py** file tried to import a module from the wrong location. Consequently, **Anaconda** failed to start and the following error message was generated:

```
No module named textw.netconfig_text
```

The code has been corrected and the error no longer occurs in the scenario described.

BZ#788537

Prior to this update, **kickstart** repository entries did not use the global proxy setting. Consequently, on networks restricted to use a proxy installation would terminate unexpectedly when attempting to connect to additional repository entries in a kickstart file if no proxy had been explicitly specified. This update changes the code to use the global proxy if an additional repository has no proxy set for it. As a result, the global proxy setting will be used and installation will proceed as expected in the scenario described.

BZ#800388

The **kickstart** pre and post installation scripts had no information about the proxy being used by **Anaconda**. As a consequence, programs such as **wget** and **curl** would not work properly in a pre-installation and post-installation script on networks restricted to using a proxy. This update sets the ***PROXY***, ***PROXY_USER***, ***PROXY_PASSWORD*** environmental variables. As a result, pre and post installation scripts now have access to the proxy setting used by Anaconda.

BZ#802397

Using the ***--onbiosdisk***=**NUMBER** option for the kickstart **part** command sometimes caused installation failures as **Anaconda** was not able to find the disk that matches the specified BIOS disk number. Users wishing to use BIOS disk numbering to control kickstart installations were not able to successfully install Red Hat Enterprise Linux. This update adjusts the comparison in Anaconda that matches the BIOS disk number to determine the Linux device name. As a result, users wishing to use BIOS disk numbering to control kickstart installations will now be able to successfully install Red Hat Enterprise Linux.

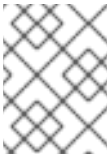
BZ#805910

Due to a regression, when running the system in Rescue mode with no or only uninitialized disks, the **Anaconda** storage subsystem did not check for the presence of a GUI before presenting the user with a list of options. Consequently, when the user selected **continue** the installer terminated

unexpectedly with a traceback. This update adds a check for presence of the GUI and falls back to a TUI if there is none. As a result, the user is informed about the lack of usable disks in the scenario described.

BZ#823810

When using **Anaconda** with Qlogic qla4xxx devices in firmware boot mode and with iSCSI targets set up in BIOS (either enabled or disabled), the devices were exposed as iSCSI devices. But in this mode the devices cannot be handled with the **iscsiadm** and **libiscsi** tools used by the installer. Consequently, installation failed with a traceback during examination of storage devices by the installer. This update changes the installer to not try to manage iSCSI devices set up with qla4xxx firmware with iscsiadm or libiscsi. As a result, installation in an environment with iSCSI targets set up by qla4xxx devices in firmware mode finishes successfully.



NOTE

The firmware boot mode is turned on and off by the ***qla4xxx.ql4xdisablesysfsboot*** boot option. With this update, it is enabled by default.

Enhancements

BZ#500273

There was no support for binding of **iSCSI** connections to network interfaces, which is required for installations using multiple iSCSI connections to a target on a single subnet for Device Mapper Multipath (DM-Multipath) connectivity. Consequently, DM-Multipath connectivity could not be used on a single subnet as all devices used the default network interface. With this update, the **Bind targets to network interfaces** option has been added to the "Advanced Storage Options" dialog box. When turned on, targets discovered specifically for all active network interfaces are available for selection and login. For kickstart installations a new ***iscsi* --iface** option can be used to specify network interface to which a target should be bound. Once interface binding is used, all iSCSI connections have to be bound, that is to say the **--iface** option has to be specified for all iscsi commands in kickstart. Network devices required for iSCSI connections can be activated either using kickstart network command with the **--activate** option or in the graphical user interface (GUI) using the **Configure Network** button from the "Advanced Storage Options" dialog ("Connect Automatically" has to be checked when configuring the device so that the device is also activated in the installer). As a result, it is now possible to configure and use DM-Multipath connectivity for iSCSI devices using different network interfaces on a single subnet during installation.

BZ#625697

The **curl** command line tool was not in the install image file. Consequently, curl could not be used in the **%pre** section of kickstart. This update adds curl to the install image and curl can be used in the **%pre** section of kickstart.

BZ#660686

Support for installation using *IP over InfiniBand* (IPoIB) interfaces has been added. As a result, it is possible to install systems connected directly to an **InfiniBand** network using IPoIB network interfaces.

BZ#663647

Two new options were added to the kickstart **volgroup** command to specify initially unused space in megabytes or as a percentage of the total volume group size. These options are only valid for volume groups being created during installation. As a result, users can effectively reserve space in a new

volume group for snapshots while still using the **--grow** option for logical volumes within the same volume group.

BZ#671230

The **GPT** disk label is now used for disks of size 2.2 TB and larger. As a result, **Anaconda** now allows installation to disks of size 2.2 TB and larger, but the installed system will not always boot properly on non-**EFI** systems. Disks of size 2.2 TB and larger may be used during the installation process, but only as data disks; they should not be used as bootable disks.

BZ#705328

When an interface configuration file is created by a configuration application such as **Anaconda**, **NetworkManager** generates the *Universally Unique IDentifier* (UUID) by hashing the existing configuration file name. Consequently, the same UUID was generated on multiple installed systems for a given network device name. With this update, a random UUID is generated by Anaconda for NetworkManager so that it does not have to generate the connection UUID by hashing the configuration file name. As a result, each network connection of all installed systems has different UUID.

BZ#735791

When **IPv6** support is set to be disabled by the installer using the **noipv6** boot option, or the ***network*** **--nopipv6** kickstart command, or by using the "Configure TCP/IP" screen of the loader *Text User Interface* (TUI), and no network device is configured for **IPv6** during installation, the IPv6 kernel modules on the installed system will now be disabled.

BZ#735857

The ability to configure a **VLAN** discovery option for *Fibre Channel over Ethernet* (FCoE) devices added during installation using **Anaconda**'s graphical user interface was required. All FCoE devices created in Anaconda installer were configured to perform VLAN discovery using the **fcoemon** daemon by setting the ***AUTO_VLAN*** value of its configuration file to **yes**. A new "Use auto vlan" checkbox was added to the "Advanced Storage Options" dialog, which is invoked by the **Add Advanced Target** button in "Advanced Storage Devices" screen. As a result, when adding FCoE device in Anaconda, it is now possible to configure the VLAN discovery option of the device using "Use auto vlan" checkbox in "Advanced Storage Options" dialog. The value of ***AUTO_VLAN*** option of FCoE device configuration file **/etc/fcoe/cfg-device** is set accordingly.

BZ#737097

The **lsscsi** and **sg3_utils** were not present in the install image. Consequently, maintenance of *Data Integrity Field* (DIF) disks was not possible. This update adds the lsscsi and sg3_utils to the install image and now utilities to maintain DIF disks can be used during the installation.

BZ#743784

Anaconda creates FCoE configuration files under the **/etc/fcoe/** directory using **biosdevname**, which is the new style interface naming scheme, for all the available Ethernet interfaces for FCoE BFS. However, it did not add the **ifname** kernel command line argument for FCoE interface that stays offline after discovering FCoE targets during installation. Because of this, during subsequent reboot the system tried to find the old style **eth*X*** interface name in **/etc/fcoe/**, which does not match the file created by Anaconda using biosdevname. Therefore, due to the missing FCoE config file, FCoE interface is never created on this interface. Consequently, during FCoE BFS installation, when an Ethernet interface went offline after discovering the targets, FCoE links did not come up after reboot. This update adds **dracut** ***ip*** parameters for all FCoE interfaces including those that went offline during installation. As a result, FCoE interfaces disconnected during installation will be activated after reboot.

BZ#744129

Installations with the ***swap* --recommended** command in kickstart created a swap file of size 2 GB plus the installed RAM size regardless of the amount of RAM installed. Consequently, machines with a large amount of RAM had huge swap files prolonging the time before the **oom_kill** syscall was invoked even in malfunctioning cases. In this update, swap size calculations for ***swap* --recommended** were changed to meet the values recommended in the documentation https://access.redhat.com/site/solutions/15244 and the **--hibernation** option was added for the **swap** kickstart command and as the default in GUI/TUI installations. As a result, machines with a lot of RAM have a reasonable swap size now if ***swap* --recommended** is used. However, hibernation might not work with this configuration. If users want to use hibernation they should use ***swap* --hibernation**.

BZ#755147

If there are multiple Ethernet interfaces configured for FCoE boot, by default, only the primary interface is turned on and the other interfaces are not configured. This update sets the value ***ONBOOT*=yes** in the **ifcfg** configuration file during installation for all network interfaces used by FCoE. As a result, all network devices used for installation to FCoE storage devices are activated automatically after reboot.

BZ#770486

This update adds the **Netcat** (**nc**) networking utility to the install environment. Users can now use the **nc** program in Rescue mode.

BZ#773545

The **virt-what** shell script has been added to the install image. Users can now use the virt-what tool in kickstart.

BZ#784327

Firmware files were loaded only from RPM files in **$prefix/lib/firmware** paths on a *Driver Update Disk* (DUD). This update adds the **$prefix/lib/firmware/updates** directory to the path to be searched for firmware. RPM files containing firmware updates can now have firmware files in **%prefix/lib/firmware/updates**.

Users of anaconda should upgrade to these updated packages, which resolve these issues and add these enhancements.

5.7. ATLAS

5.7.1. RHBA-2012:0402 – atlas bug fix update

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

The ATLAS (Automatically Tuned Linear Algebra Software) project is a research effort focusing on applying empirical techniques providing portable performance. The atlas packages provide C and Fortran77 interfaces to a portably efficient BLAS (Basic Linear Algebra Subprograms) implementation and routines from LAPACK (Linear Algebra PACKKage).

Bug Fix

BZ#723350

Previously, binary files from the base atlas package contained illegal instructions from an incompatible instruction set (3DNow!). As a consequence, an "Illegal instruction" error was displayed. This update disables usage of the instruction set.

All users of atlas are advised to upgrade to these updated packages, which fix this bug.

5.8. AUDIT

5.8.1. RHBA-2012:0929 – audit bug fix and enhancement update

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

The audit packages contain the user space utilities for storing and searching the audit records which have been generated by the audit subsystem in the Linux 2.6 kernel.

The audit packages have been upgraded to upstream version 2.2, which provides a number of bug fixes and enhancements over the previous version. The version 2.2 packages introduce the following enhancements:

- The "auditctl" command now allows shell-escaped file names for better handling of file names with spaces in them.
- There is a new utility, auvirt, that extracts a report about the virtualization events.
- The auditd.conf configuration option, "tcp_max_per_addr", now allows up to 1024 concurrent connections from the same IP address. While this is not recommended for normal use, it helps in situations where a number of client systems are behind a NAT, which causes them to appear to have the same IP address.

Bug Fixes

BZ#803349

Previously, not enough information was parsed to determine whether audit records are part of the same event if the server's node name was longer than approximately 80 characters. With this update, the problem has been fixed.

BZ#797848

This update fixes a typo in the audit.rules(7) man page.

Enhancements

BZ#658630

Prior to this update, if the audit rules had a typo or the command was not supported by the Linux kernel, either an error was triggered and you were able to stop processing the rules or, as the other option, you were able to ignore any errors in which case it completed everything it could but returned success. This update introduces the "-c" option to auditctl which works like the ignore option, but instead of returning success, the "-c" option returns failure if any rule triggers an error. Note that like the ignore option, the "-c" option continues to process all audit rules.

BZ#766920

This release adds support for a new kernel auditing feature that allows for inter-field comparisons.

For each audit event, the Linux kernel collects information about what is causing the event. Now, you can use the "-C" option to compare: "auid", "uid", "euid", "suid", "fsuid", or "obj_uid"; and "gid", "egid", "sgid", "fsgid", or "obj_gid". The two groups cannot be mixed. Comparisons can use either the equal or not equal operators. Note that for this enhancement to work, the system must boot the Linux 2.6.32-244 kernel or later.

All audit users are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.9. AUGEAS

5.9.1. RHBA-2012:0967 – augeas bug fix and enhancement update

Updated augeas packages that fix three bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

Augeas is a configuration editing tool. 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.

Bug Fixes

BZ#759311

Previously, the "--autosave" option did not work correctly when using Augeas in batch mode, which caused that configuration changes were not saved. As a consequence, configuration changes could be saved only in interactive mode. This update ensures that the "--autosave" option functions in batch mode as expected.

BZ#781690

Prior to this update, when parsing GRUB configuration files, Augeas did not parse the "--encrypted" option of the "password" command correctly. Instead, it parsed the "--encrypted" part as the password, and the password hash as a second "menu.lst" filename. This update ensures that the "--encrypted" option of the password command is parsed correctly when parsing GRUB configuration files.

BZ#820864

Previously, Augeas was not able to parse the /etc/fstab file containing mount options with an equals sign but no value. This update fixes the fstab lens so that it can handle such mount options. As a result, Augeas can now parse an /etc/fstab file containing mount options with an equals sign but no value correctly.

Enhancements

BZ#628507

Previously, the finite-automata-DOT graph tool (fadot) did not support the -h option. Consequently, when fadot was launched with the -h option the "Unknown option" message was displayed. This update adds support for the -h option and ensures that a help message is displayed when fadot is launched with the option.

BZ#808662

Previously, Augeas did not have a lens to parse the /etc/mdadm.conf file. Consequently, the tool for

conversion of physical servers to virtual guests, Virt-P2V, could not convert physical hosts on MD devices. This update adds a new lens to parse the /etc/mdadm.conf file, enabling Virt-P2V to convert physical hosts on MD devices as expected.

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

5.10. AUTHCONFIG

5.10.1. RHBA-2012:0931 – authconfig bug fix and enhancement update

Updated authconfig packages that fix multiple bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The authconfig packages provide a command line utility and a GUI application that can configure a workstation to be a client for certain network user information and authentication schemes, and other user information and authentication related options.

Bug Fixes

BZ#689717

Prior to this update, SSSD configuration files failed to parse if the files were not correctly formatted. As a consequence, the authconfig utility could abort unexpectedly. With this update, the error is correctly handled, the configuration file is backed up, and a new file is created.

BZ#708850

Prior to this update, the man page "authconfig(8)" referred to non-existing obsolete configuration files. This update modifies the man page to point to configuration files that are currently modified by authconfig.

BZ#749700

Prior to this update, a deprecated "krb_kdcip" option was set instead of the "krb5_server" option when the SSSD configuration was updated. This update modifies the SSSD configuration setting to use the "krb5_server" option to set the Kerberos KDC server address.

BZ#755975

Prior to this update, the authconfig command always returned the exit value "1" when the "--savebackup" option was used, due to handling of nonexisting configuration files on the system. With this update, the exit value is "0" if the configuration backup succeeds even if some configuration files which can be handled by authconfig, are not present on the system.

Enhancements

BZ#731094

Prior to this update, the authconfig utility did not support the SSSD configuration with the IPA backend. This update allows to join an IPAv2 domain with the system via the ipa-client-install command.

BZ#804615

With this update, the nss_sss module is also used in the "services" entry of the nsswitch.conf file when configuring this file.

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

5.11. AUTOFS

5.11.1. RHBA-2012:1442 – autofs bug fix update

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

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

Bug Fix

BZ#870929

During the boot-up sequence, when the automount daemon was using an internal host map, automount terminated unexpectedly with a segmentation fault. This bug has been fixed and the crashes no longer occur in the described scenario.

All users of autofs are advised to upgrade to these updated packages, which fix this bug.

5.11.2. RHBA-2012:0951 – autofs bug fix and enhancement update

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

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

Bug Fixes

BZ#772946

A recent change to correct a problem with included map entry removal introduced a new problem with included map key look-up. The condition used in the previous patch was too broad and the map key lookup mechanism failed to find keys in an included multi-mount map entry. The condition has been modified so that keys in multi-mount map entries are now found correctly.

BZ#772356

A function that checks validity of a mount location was meant to check only for a small subset of map location errors. A recent improvement modification in error reporting inverted a logic test in this validating function. Consequently, the scope of the test was widened, which caused automount to report false-positive failures. With this update, the faulty logic test has been fixed and false-positive failures no longer occur.

BZ#790674

Previously, autofs submounts incorrectly handled shutdown synchronization and lock restrictions. As a consequence, automount could become unresponsive when submounts expired. With this update, the submount shuts down only after passing through the state ST_SHUTDOWN,

ST_SHUTDOWN_PENDING, or ST_SHUTDOWN_FORCE, or when the state changes to ST_READY.

BZ#753964

Prior to this update, two IPv6 compatibility functions were erroneously not included in the autofs interface to the libtirpc library. This prevented the autofs IPv6 RPC code from working. With this update, the libtirpc interface code for autofs has been fixed.

BZ#782169

When using the legacy auto.net script for the hosts map, an error in the script for handling multiple occurrences of exports prevented the script from returning any of the exported paths. This bug has been fixed by modifying the script to select only a unique list of exports, thus eliminating duplicate exports.

BZ#787595

Due to changes to the mount.nfs utility to take advantage of the support for NFS mount options in the kernel, the RPC processing had moved from mount.nfs to the kernel. However, the kernel RPC had to wait for RPC requests to servers that were not available to time out, resulting in very slow interactive response when attempting an automount to a server that was not available. This update changes the autofs RPC code to detect this situation early and provide proper error messages as soon as possible.

BZ#760945

Previously, although the /net/ and /misc/ directories are exclusively used by the default /etc/auto.master utility, they were not specified in the autofs RPM package. As a result, the rpm utility reported them as not owned by any package. This update adds both these directories to the autofs spec file.

BZ#745527

Previously, the autofs init.d script failed to return proper usage messages if called with no arguments, or incorrect arguments. This bug has been fixed and the script now prints the usage information as expected.

Enhancement

BZ#683523

Initial support for the System Security Services Daemon (SSSD) as a map source has been added to the autofs package.

All autofs users are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.12. AXIS

5.12.1. RHSA-2013:0269 – Moderate: axis security update

Updated axis packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

Apache Axis is an implementation of SOAP (Simple Object Access Protocol). It can be used to build both web service clients and servers.

Security Fix

CVE-2012-5784

Apache Axis did not verify that the server hostname matched the domain name in the subject's Common Name (CN) or subjectAltName field in X.509 certificates. This could allow a man-in-the-middle attacker to spoof an SSL server if they had a certificate that was valid for any domain name.

All users of axis are advised to upgrade to these updated packages, which correct this issue. Applications using Apache Axis must be restarted for this update to take effect.

5.13. BACULA

5.13.1. RHBA-2012:1469 – bacula bug fix update

Updated bacula packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The bacula packages provide a tool set that allows you to manage the backup, recovery, and verification of computer data across a network of different computers.

Bug Fixes

BZ#728693

Prior to this update, the logwatch tool did not check the "/var/log/bacula*" file. As a consequence, the logwatch report was incomplete. This update adds all log files to the logwatch configuration file. Now, the logwatch report is complete.

BZ#728697

Prior to this update, the bacula tool itself created the "/var/spool/bacula/log" file. As a consequence, this log file used an incorrect SELinux context. This update modifies the underlying code to create the /var/spool/bacula/log file in the bacula package. Now, this log file has the correct SELinux context.

BZ#729008

Prior to this update, the bacula packages were built without the CFLAGS variable "$RPM_OPT_FLAGS". As a consequence, the debug information was not generated. This update modifies the underlying code to build the packages with CFLAGS="$RPM_OPT_FLAGS. Now, the debug information is generated as expected.

BZ#756803

Prior to this update, the perl script which generates the my.conf file contained a misprint. As a consequence, the port variable was not set correctly. This update corrects the misprint. Now, the port variable is set as expected.

BZ#802158

Prior to this update, values for the "show pool" command was obtained from the "res->res_client" item. As a consequence, the output displayed incorrect job and file retention values. This update uses the "res->res_pool" item to obtain the correct values.

BZ#862240

Prior to this update, bacula-storage-common utility wrongly removed alternatives for the bcopy function during the update. As a consequence, the Link to bcop.{mysql,sqlite,postgresql} disappeared after updating. This update modifies the underlying code to remove these links directly in storage-{mysql,sqlite,postgresql} and not in bacula-storage-common.

All users of bacula are advised to upgrade to these updated packages, which fix these bugs.

5.14. BIND-DYNDB-LDAP

5.14.1. RHSA-2012:1139 – Important: bind-dyndb-ldap security update

An updated bind-dyndb-ldap package that fixes one security issue is now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

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 help to reduce the load on LDAP servers.

Security Fix

CVE-2012-3429

A flaw was found in the way bind-dyndb-ldap performed the escaping of names from DNS requests for use in LDAP queries. A remote attacker able to send DNS queries to a named server that is configured to use bind-dyndb-ldap could use this flaw to cause named to exit unexpectedly with an assertion failure.

Red Hat would like to thank Sigbjorn Lie of Atea Norway for reporting this issue.

All bind-dyndb-ldap users should upgrade to this updated package, which contains a backported patch to correct this issue. For the update to take effect, the named service must be restarted.

5.14.2. RHBA-2012:0837 – bind-dyndb-ldap bug fix and enhancement update

An updated bind-dyndb-ldap package which provides a number of bug fixes and enhancements is 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 help to reduce the load on LDAP servers.



NOTE

The bind-dyndb-ldap package has been upgraded to upstream version **1.1.0b2**, which provides a number of bug fixes and enhancements over the previous version (BZ#767486).

Bug Fixes

BZ#751776

The bind-dyndb-ldap plug-in refused to load an entire zone when it contained an invalid *Resource Record* (RR) with the same *Fully Qualified Domain Name* (FQDN) as the zone name (for example an MX record). With this update, the code for parsing Resource Records has been improved. If an invalid RR is encountered, an error message "Failed to parse RR entry " is logged and the zone continues to load successfully.

BZ#767489

When the first connection to an **LDAP** server failed, the bind-dyndb-ldap plug-in did not try to connect again. Consequently, users had to execute the "rndc reload" command to make the plug-in work. With this update, the plug-in periodically retries to connect to an LDAP server. As a result, user intervention is no longer required and the plug-in works as expected.

BZ#767492

When the ***zone_refresh*** period timed out and a zone was removed from the **LDAP** server, the plug-in continued to serve the removed zone. With this update, the plug-in no longer serves zones which have been deleted from LDAP when the ***zone_refresh*** parameter is set.

BZ#789356

When the named daemon received the **rndc reload** command or a **SIGHUP** signal and the plug-in failed to connect to an LDAP server, the plug-in caused named to terminate unexpectedly when it received a query which belonged to a zone previously handled by the plug-in. This has been fixed, the plug-in no longer serves its zones when connection to LDAP fails during reload and no longer crashes in the scenario described.

BZ#796206

The plug-in terminated unexpectedly when named lost connection to an **LDAP** server for some time, then reconnected successfully, and some zones previously present had been removed from the LDAP server. The bug has been fixed and the plug-in no longer crashes in the scenario described.

BZ#805871

Certain string lengths were incorrectly set in the plug-in. Consequently, the *Start of Authority* (SOA) serial number and expiry time were incorrectly set for the forward zone during **ipa-server** installation. With this update, the code has been improved and the SOA serial number and expiry time are set as expected.

BZ#811074

When a *Domain Name System* (DNS) zone was managed by a bind-dyndb-ldap plugin and a sub-domain was delegated to another **DNS** server, the plug-in did not put A or AAAA glue records in the "additional section" of a DNS answer. Consequently, the delegated sub-domain was not accessible by other DNS servers. With this update, the plug-in has been fixed and now returns A or AAAA glue records of a delegated sub-domain in the "additional section". As a result, delegated zones are correctly resolvable in the scenario described.

BZ#818933

Previously, the bind-dyndb-ldap plug-in did not escape non-ASCII characters in incoming DNS queries correctly. Consequently, the plug-in failed to send answers for queries which contained non-ASCII characters such as ",". The plug-in has been fixed and now correctly returns answers for queries with non-ASCII characters.

Enhancements

BZ#733371

The bind-dyndb-ldap plug-in now supports two new attributes, ***idnsAllowQuery*** and ***idnsAllowTransfer***, which can be used to set ACLs for queries or transfers. Refer to **/usr/share/doc/bind-dyndb-ldap/README** for information on the attributes.

BZ#754433

The plug-in now supports the new zone attributes ***idnsForwarders*** and ***idnsForwardPolicy*** which can be used to configure forwarding. Refer to **/usr/share/doc/bind-dyndb-ldap/README** for a detailed description.

BZ#766233

The plug-in now supports zone transfers.

BZ#767494

The plug-in has a new option called ***sync_ptr*** that can be used to keep A and AAAA records and their PTR records synchronized. Refer to **/usr/share/doc/bind-dyndb-ldap/README** for a detailed description.

BZ#795406

It was not possible to store configuration for the plug-in in **LDAP** and configuration was only taken from the **named.conf** file. With this update, configuration information can be obtained from ***idnsConfigObject*** in LDAP. Note that options set in named.conf have lower priority than options set in LDAP. The priority will change in future updates. Refer to the README file for more details.

Users of bind-dyndb-ldap package should upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.15. BIND

5.15.1. RHBA-2012:1107 – bind bug fix update

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

BIND (Berkeley Internet Name Domain) is an implementation of the DNS (Domain Name System) protocols. BIND includes a DNS server (named), which resolves host names to IP addresses; a resolver library (routines for applications to use when interfacing with the DNS server); and tools for verifying that the DNS server is operating properly.

Bug Fix

BZ#838956

Due to a race condition in the rbtdb.c source file, the named daemon could terminate unexpectedly with the INSIST error code. This bug has been fixed in the code and the named daemon no longer crashes in the described scenario.

All users of bind are advised to upgrade to these updated packages, which fix this bug.

5.15.2. RHSA-2012:1549 – Important: bind security update

Updated bind packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

The Berkeley Internet Name Domain (BIND) is an implementation of the Domain Name System (DNS) protocols. BIND includes a DNS server (named); a resolver library (routines for applications to use when interfacing with DNS); and tools for verifying that the DNS server is operating correctly. DNS64 is used to automatically generate DNS records so IPv6 based clients can access IPv4 systems through a NAT64 server.

Security Fix

CVE-2012-5688

A flaw was found in the DNS64 implementation in BIND. If a remote attacker sent a specially-crafted query to a named server, named could exit unexpectedly with an assertion failure. Note that DNS64 support is not enabled by default.

Users of bind are advised to upgrade to these updated packages, which correct this issue. After installing the update, the BIND daemon (named) will be restarted automatically.

5.15.3. RHSA-2012:1268 – Important: bind security update

Updated bind packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

The Berkeley Internet Name Domain (BIND) is an implementation of the Domain Name System (DNS) protocols. BIND includes a DNS server (named); a resolver library (routines for applications to use when interfacing with DNS); and tools for verifying that the DNS server is operating correctly.

Security Fix

CVE-2012-4244

A flaw was found in the way BIND handled resource records with a large RDATA value. A malicious owner of a DNS domain could use this flaw to create specially-crafted DNS resource records, that would cause a recursive resolver or secondary server to exit unexpectedly with an assertion failure.

Users of bind are advised to upgrade to these updated packages, which correct this issue. After installing the update, the BIND daemon (named) will be restarted automatically.

5.15.4. RHSA-2012:1123 – Important: bind security update

Updated bind packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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(s) associated with each description below.

The Berkeley Internet Name Domain (BIND) is an implementation of the Domain Name System (DNS) protocols. BIND includes a DNS server (named); a resolver library (routines for applications to use when interfacing with DNS); and tools for verifying that the DNS server is operating correctly.

Security Fix

CVE-2012-3817

An uninitialized data structure use flaw was found in BIND when DNSSEC validation was enabled. A remote attacker able to send a large number of queries to a DNSSEC validating BIND resolver could use this flaw to cause it to exit unexpectedly with an assertion failure.

Users of bind are advised to upgrade to these updated packages, which correct this issue. After installing the update, the BIND daemon (named) will be restarted automatically.

5.15.5. RHBA-2012:1341 – bind bug fix update

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

BIND (Berkeley Internet Name Domain) is an implementation of the DNS (Domain Name System) protocols. BIND includes a DNS server (named), which resolves host names to IP addresses; a resolver library containing routines for applications to use when interfacing with the DNS server; and tools for verifying that the DNS server is operating properly.

Bug Fix

BZ#858273

Previously, BIND rejected "forward" and "forwarders" statements in static-stub zones. Consequently, it was impossible to forward certain queries to specified servers. With this update, BIND accepts those options for static-stub zones properly, thus fixing this bug.

All users of bind are advised to upgrade to these updated packages, which fix this bug.

5.15.6. RHSA-2012:1363 – Important: bind security update

Updated bind packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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(s) associated with each description below.

The Berkeley Internet Name Domain (BIND) is an implementation of the Domain Name System (DNS) protocols. BIND includes a DNS server (named); a resolver library (routines for applications to use when interfacing with DNS); and tools for verifying that the DNS server is operating correctly.

Security Fix

CVE-2012-5166

A flaw was found in the way BIND handled certain combinations of resource records. A remote attacker could use this flaw to cause a recursive resolver, or an authoritative server in certain configurations, to lockup.

Users of bind are advised to upgrade to these updated packages, which correct this issue. After installing the update, the BIND daemon (named) will be restarted automatically.

5.15.7. RHBA-2012:0830 – bind bug fix and enhancement update

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

BIND (*Berkeley Internet Name Domain*) is an implementation of the **DNS** (*Domain Name System*) protocols. BIND includes a DNS server (**named**), which resolves host names to IP addresses; a resolver library (routines for applications to use when interfacing with DNS); and tools for verifying that the DNS server is operating properly.



NOTE

The bind package has been upgraded to upstream version **9.8.2rc1** which provides a number of bug fixes and enhancements over the previous version. Refer to **/usr/share/doc/bind-9.8.2/README** for a detailed list of enhancements. (BZ# 745284, BZ#755618, BZ#797972)

Bug Fixes

BZ#734458

When **/etc/resolv.conf** contained nameservers with disabled recursion, **nslookup** failed to resolve certain host names. With this update, a patch has been applied and **nslookup** now works as expected in the scenario described.

BZ#739406

Prior to this update, errors arising on automatic update of **DNSSEC** trust anchors were handled incorrectly. Consequently, the **named** daemon could become unresponsive on shutdown. With this update, the error handling has been improved and **named** exits on shutdown gracefully.

BZ#739410

The multi-threaded **named** daemon uses the atomic operations feature to speed-up access to shared data. This feature did not work correctly on 32-bit and 64-bit PowerPC architectures. Therefore, **named** sometimes became unresponsive on these architectures. This update disables the atomic operations feature on 32-bit and 64-bit PowerPC architectures, which ensures that **named** is now more stable and reliable and no longer hangs.

BZ#746694

Prior to this update, a race condition could occur on validation of DNSSEC-signed NXDOMAIN responses and **named** could terminate unexpectedly. With this update, the underlying code has been fixed and the race condition no longer occurs.

BZ#759502

The **named** daemon, configured as the master server, sometimes failed to transfer an uncompressible zone. The following error message was logged:

```
transfer of './IN': sending zone data: ran out of space
```

The code which handles zone transfers has been fixed and this error no longer occurs in the scenario described.

BZ#759503

During a **DNS** zone transfer, **named** sometimes terminated unexpectedly with an assertion failure. With this update, a patch has been applied to make the code more robust, and **named** no longer crashes in the scenario described.

BZ#768798

Previously, the **rndc.key** file was generated during package installation by the **rndc-confgen -a** command, but this feature was removed in Red Hat Enterprise Linux 6.1 because users reported that installation of bind package sometimes hung due to lack of entropy in **/dev/random**. The **named** initscript now generates **rndc.key** during the service startup if it does not exist.

BZ#786362

After the **rndc reload** command was executed, **named** failed to update **DNSSEC** trust anchors and emitted the following message to the log:

```
managed-keys-zone ./IN: Failed to create fetch for DNSKEY update
```

This issue was fixed in the 9.8.2rc1 upstream version.

BZ#789886

Due to an error in the bind spec file, the bind-chroot subpackage did not create a **/dev/null** device. In addition, some empty directories were left behind after uninstalling bind. With this update, the bind-chroot packaging errors have been fixed.

BZ#795414

The dynamic-db plug-ins were loaded too early which caused the configuration in the **named.conf** file to override the configuration supplied by the plug-in. Consequently, **named** sometimes failed to start. With this update the **named.conf** is parsed before plug-in initialization and **named** now starts as expected.

BZ#812900

Previously, when the **/var/named** directory was mounted the **/etc/init.d/named** initscript did not distinguish between situations when ***chroot*** configuration was enabled and when ***chroot*** was not enabled. Consequently, when stopping the **named** service the **/var/named** directory was always unmounted. The initscript has been fixed and now unmounts **/var/named** only when ***chroot*** configuration is enabled. As a result, **/var/named** stays mounted after the **named** service is stopped when ***chroot*** configuration is not enabled.

BZ#816164

Previously, the **nslookup** utility did not return a non-zero exit code when it failed to get an answer. Consequently, it was impossible to determine if an nslookup run was successful or not from the error code. The nslookup utility has been fixed and now it returns "1" as the exit code when fails to get answer.

Enhancements

BZ#735438

By default **BIND** returns resource records in round-robin order. The ***rrset-order*** option now supports **fixed** ordering. When this option is set, the resource records for each domain name are always returned in the order they are loaded from the zone file.

BZ#788870

Previously, **named** logged too many messages relating to external **DNS** queries. The severity of these error messages has been decreased from "notice" to "debug" so that the system log is not flooded with mostly unnecessary information.

BZ#790682

The **named** daemon now uses **portreserve** to reserve the *Remote Name Daemon Control* (RNDC) port to avoid conflicts with other services.

All users of bind are advised to upgrade to these updated packages, which fix these bugs and provide these enhancements.

5.16. BINUTILS

5.16.1. RHBA-2012:0872 – binutils bug fix and enhancement update

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

The binutils packages contain a collection of binary utilities, including "ar" (for creating, modifying and extracting from archives), "as" (a family of GNU assemblers), "gprof" (for displaying call graph profile data), "ld" (the GNU linker), "nm" (for listing symbols from object files), "objcopy" (for copying and translating object files), "objdump" (for displaying information from object files), "ranlib" (for generating an index for the contents of an archive), "readelf" (for displaying detailed information about binary files), "size" (for listing the section sizes of an object or archive file), "strings" (for listing printable strings from files), "strip" (for discarding symbols), and "addr2line" (for converting addresses to file and line).

Bug Fixes

BZ#676194

Previously, the GNU linker could terminate unexpectedly with a segmentation fault when attempting to link together object files of different architectures (for example, an object file of 32-bit Intel P6 with an object file of Intel 64). This update modifies binutils so that the linker now generates an error message and refuses to link object files in the scenario described.

BZ#809616

When generating build-ID hashes, the GNU linker previously allocated memory for BSS sections. Consequently, the linker could use more memory than was necessary. This update modifies the linker to skip BSS sections and thus avoid unnecessary memory usage when generating build-ID hashes.

Enhancements

BZ#739444

With this update, backported patches have been included to support new AMD processors. Also, a duplicate entry for the bextr instruction has been removed from the disassembler's table.

BZ#739144

The GNU linker has been modified in order to improve performance of Table of Contents (TOC) addressability and Procedure Linkage Table (PLT) call stubs on the PowerPC and PowerPC 64 architectures.

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

5.17. BIOSDEVNAME

5.17.1. RHBA-2013:0138 – biosdevname bug fix update

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

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

Bug Fix

BZ#865446

Previously, biosdevname did not handle PCI cards with multiple ports properly. Consequently, only the network interface of the first port of these cards was renamed according to the biosdevname naming scheme. This bug has been fixed and network interfaces of all ports of these cards are now renamed as expected.

Users of biosdevname are advised to upgrade to these update packages, which fix this bug.

5.18. BRLTTY

5.18.1. RHBA-2012:1231 – brltty bug fix update

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

BRLTTY is a background process (daemon) which provides access to the Linux console (when in text mode) for a blind person using a refreshable braille display. It drives the braille display, and provides complete screen review functionality.

Bug Fixes

BZ#684526

Previously, building the brltty package could fail on the ocaml's unpackaged files error. This happened only if the ocaml package was pre-installed in the build root. The "--disable-caml-bindings" option has been added in the %configure macro so that the package now builds correctly.

BZ#809326

Previously, the /usr/lib/libbrlapi.so symbolic link installed by the brlapi-devel package incorrectly pointed to ../../lib/libbrlapi.so. The link has been fixed to correctly point to ../../lib/libbrlapi.so.0.5.

All users of brltty are advised to upgrade to these updated packages, which fix these bugs.

5.19. BUSYBOX

5.19.1. RHSA-2012:0810 — Low: busybox security and bug fix update

Updated busybox packages that fix two 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 low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each 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 Fixes

CVE-2006-1168

A buffer underflow flaw was found in the way the uncompress utility of BusyBox expanded certain archive files compressed using Lempel-Ziv compression. If a user were tricked into expanding a specially-crafted archive file with uncompress, it could cause BusyBox to crash or, potentially, execute arbitrary code with the privileges of the user running BusyBox.

CVE-2011-2716

The BusyBox DHCP client, udhcpc, did not sufficiently sanitize certain options provided in DHCP server replies, such as the client hostname. A malicious DHCP server could send such an option with a specially-crafted value to a DHCP client. If this option's value was saved on the client system, and then later insecurely evaluated by a process that assumes the option is trusted, it could lead to arbitrary code execution with the privileges of that process. Note: udhcpc is not used on Red Hat Enterprise Linux by default, and no DHCP client script is provided with the busybox packages.

Bug Fixes

BZ#751927

Prior to this update, the "findfs" command did not recognize Btrfs partitions. As a consequence, an error message could occur when dumping a core file. This update adds support for recognizing such partitions so the problem no longer occurs.

BZ#752134

If the "grep" command was used with the "-F" and "-i" options at the same time, the "-i" option was ignored. As a consequence, the "grep -iF" command incorrectly performed a case-sensitive search instead of an insensitive search. A patch has been applied to ensure that the combination of the "-F" and "-i" options works as expected.

BZ#782018

Prior to this update, the msh shell did not support the "set -o pipefail" command. This update adds support for this command.

BZ#809092

Previously, the msh shell could terminate unexpectedly with a segmentation fault when attempting to execute an empty command as a result of variable substitution (for example msh -c '$nonexistent_variable'). With this update, msh has been modified to correctly interpret such commands and no longer crashes in this scenario.

BZ#752132

Previously, the msh shell incorrectly executed empty loops. As a consequence, msh never exited such a loop even if the loop condition was false, which could cause scripts using the loop to become unresponsive. With this update, msh has been modified to execute and exit empty loops correctly, so that hangs no longer occur.

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

5.20. BYACC

5.20.1. RHBA-2012:0749 – byacc bug fix update

An updated byacc package that fixes one bug is now available for Red Hat Enterprise Linux 6.

Berkeley Yacc (byacc) is a public domain look-ahead left-to-right (LALR) parser generator used by many programs during their build process.

Bug Fix

BZ#743343

Byacc's maximum stack depth was reduced from 10000 to 500 between byacc releases. If deep enough else-if structures were present in source code being compiled with byacc, this could lead to out-of-memory conditions, resulting in YACC Stack Overflow and build failure. This updated release restores the maximum stack depth to its original value, 10000. Note: the underlying LR algorithm still imposes a hard limit on the number of parsable else-if statements. Restoring the maximum stack depth to its original value means source code with deep else-if structures that previously compiled against byacc will again do so.

All byacc users should upgrade to this updated package, which fixes this bug.

5.21. C-ARES

5.21.1. RHBA-2012:0922 – c-ares bug fix update

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

The c-ares C library defines asynchronous DNS (Domain Name System) requests and provides name resolving API.

Bug Fixes

BZ#730695

Previously, when searching for AF_UNSPEC or AF_INET6 address families, the c-ares library fell back to the AF_INET family if no AF_INET6 addresses were found. Consequently, IPv4 addresses were returned even if only IPv6 addresses were requested. With this update, c-ares performs the

fallback only when searching for AF_UNSPEC addresses.

BZ#730693

The ares_parse_a_reply() function leaked memory when the user attempted to parse an invalid reply. With this update, the allocated memory is freed properly and the memory leak no longer occurs.

BZ#713133

A switch statement inside the ares_malloc_data() public function was missing a terminating break statement. This could result in unpredictable behavior and sometimes the application terminated unexpectedly. This update adds the missing switch statement and the ares_malloc_data() function now works as intended.

BZ#695426

When parsing SeRVice (SRV) record queries, c-ares was accessing memory incorrectly on architectures that require data to be aligned in memory. This caused the program to terminate unexpectedly with the SIGBUS signal. With this update, c-ares has been modified to access the memory correctly in the scenario described.

BZ#640944

Previously, the ares_gethostbyname manual page did not document the ARES_ENODATA error code as a valid and expected error code. With this update, the manual page has been modified accordingly.

All users of c-ares are advised to upgrade to these updated packages, which fix these bugs.

5.22. CDRKIT

5.22.1. RHBA-2012:1451 – cdrkit bug fix update

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

The cdrkit packages contain a collection of CD/DVD utilities for generating the ISO9660 file-system and burning media.

Bug Fix

BZ#797990

Prior to this update, overlapping memory was handled incorrectly. As a consequence, newly created paths could be garbled when calling "genisoimage" with the "-graft-points" option to graft the paths at points other than the root directory. This update modifies the underlying code to generate graft paths as expected.

All users of cdrkit are advised to upgrade to these updated packages, which fix this bug.

5.23. CERTMONGER

5.23.1. RHBA-2012:0833 – certmonger bug fix and enhancement update

Updated certmonger packages that fix multiple bugs and add multiple enhancements are now available for Red Hat Enterprise Linux 6.

The certmonger daemon monitors certificates which have been registered with it, and as a certificate's not-valid-after date approaches, the daemon can optionally attempt to obtain a fresh certificate from a supported CA.

The certmonger packages have been upgraded to upstream version 0.56, which provides a number of bug fixes and enhancements over the previous version. (BZ#789153)

Bug Fixes

BZ#765599

Prior to this update, one of the examples provided in the getting-started.txt file did not work as expected if the daemon was prevented from accessing files in user-specified locations, for example by the SELinux policy. With this update, this problem is now documented in the getting-started.txt file.

BZ#765600

Prior to this update, the certmonger daemon was not configured to start by default when the package was installed. This update enables the certmonger service by default.

BZ#796542

Prior to this update, the "getcert" command could under certain circumstances, display the misleading error message "invalid option" when an option that required an argument was used and the argument was not specified. This update modifies the error code so that the correct message is now sent.

Enhancement

BZ#766167

Prior to this update, newly added certificates were not automatically visible. To see these certificates, servers had to be manually restarted. This update adds the emission of D-Bus signals over the message bus to allow applications to perform the actions they need to use a new certificate. Also, the new "-C" option was added to invoke a user-specified command.

All users of certmonger are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.24. CHKCONFIG

5.24.1. RHBA-2012:0873 – chkconfig bug fix update

Updated chkconfig packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The basic system utility chkconfig updates and queries runlevel information for system services.

Bug Fixes

BZ#696305

When installing multiple Linux Standard Base (LSB) services which only had LSB headers, the stop priority of the related LSB init scripts could have been miscalculated and set to "-1". With this update,

the LSB init script ordering mechanism has been fixed, and the stop priority of the LSB init scripts is now set correctly.

BZ#706854

When an LSB init script requiring the "$local_fs" facility was installed with the "install_initd" command, the installation of the script could fail under certain circumstances. With this update, the underlying code has been modified to ignore this requirement because the "$local_fs" facility is always implicitly provided. LSB init scripts with requirements on "$local_fs" are now installed correctly.

BZ#771454

If an LSB init script contained "Required-Start" dependencies, but the LSB service installed was not configured to start in any runlevel, the dependencies could have been applied incorrectly. Consequently, the installation of the LSB service failed silently. With this update, chkconfig no longer strictly enforces "Required-Start" dependencies for installation if the service is not configured to start in any runlevel. LSB services are now installed as expected in this scenario.

BZ#771741

Previously, chkconfig did not handle dependencies between LSB init scripts correctly. Therefore, if an LSB service was enabled, LSB services that were depending on it could have been set up incorrectly. With this update, chkconfig has been modified to determine dependencies properly, and dependent LSB services are now set up as expected in this scenario.

All users of chkconfig are advised to upgrade to these updated packages, which fix these bugs.

5.25. CIFS-UTILS

5.25.1. RHSA-2012:0902 – Low: cifs-utils security, bug fix, and enhancement update

An updated cifs-utils package that fixes one security issue, multiple bugs, and adds various 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

The cifs-utils package contains tools for mounting and managing shares on Linux using the SMB/CIFS protocol. The CIFS shares can be used as standard Linux file systems.

Security Fix

CVE-2012-1586

A file existence disclosure flaw was found in mount.cifs. If the tool was installed with the setuid bit set, a local attacker could use this flaw to determine the existence of files or directories in directories not accessible to the attacker.



NOTE

mount.cifs from the cifs-utils package distributed by Red Hat does not have the setuid bit set. We recommend that administrators do not manually set the setuid bit for mount.cifs.

Bug Fixes

BZ#769923

The cifs.mount(8) manual page was previously missing documentation for several mount options. With this update, the missing entries have been added to the manual page.

BZ#770004

Previously, the mount.cifs utility did not properly update the "/etc/mtab" system information file when remounting an existing CIFS mount. Consequently, mount.cifs created a duplicate entry of the existing mount entry. This update adds the del_mtab() function to cifs.mount, which ensures that the old mount entry is removed from "/etc/mtab" before adding the updated mount entry.

BZ#796463

The mount.cifs utility did not properly convert user and group names to numeric UIDs and GIDs. Therefore, when the "uid", "gid" or "cruid" mount options were specified with user or group names, CIFS shares were mounted with default values. This caused shares to be inaccessible to the intended users because UID and GID is set to "0" by default. With this update, user and group names are properly converted so that CIFS shares are now mounted with specified user and group ownership as expected.

BZ#805490

The cifs.upcall utility did not respect the "domain_realm" section in the "krb5.conf" file and worked only with the default domain. Consequently, an attempt to mount a CIFS share from a different than the default domain failed with the following error message:

mount error(126): Required key not available

This update modifies the underlying code so that cifs.upcall handles multiple Kerberos domains correctly and CIFS shares can now be mounted as expected in a multi-domain environment.

Enhancements

BZ#748756

The cifs.upcall utility previously always used the "/etc/krb5.conf" file regardless of whether the user had specified a custom Kerberos configuration file. This update adds the "--krb5conf" option to cifs.upcall allowing the administrator to specify an alternate krb5.conf file. For more information on this option, refer to the cifs.upcall(8) manual page.

BZ#748757

The cifs.upcall utility did not optimally determine the correct service principal name (SPN) used for Kerberos authentication, which occasionally caused krb5 authentication to fail when mounting a server's unqualified domain name. This update improves cifs.upcall so that the method used to determine the SPN is now more versatile.

BZ#806337

This update adds the "backupuid" and "backupgid" mount options to the mount.cifs utility. When specified, these options grant a user or a group the right to access files with the backup intent. For more information on these options, refer to the mount.cifs(8) manual page.

All users of cifs-utils are advised to upgrade to this updated package, which contains backported patches to fix these issues and add these enhancements.

5.26. CLUSTER AND GFS2-UTILS

5.26.1. RHBA-2012:0861 – cluster and gfs2-utils bug fix and enhancement update

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

The cluster and gfs2-utils packages contain the core clustering libraries for Red Hat High Availability as well as utilities to maintain GFS2 file systems for users of Red Hat Resilient Storage.

Bug Fixes

BZ#759603

A race condition existed when a node lost contact with the quorum device at the same time as the token timeout period expired. The nodes raced to fence, which could lead to a cluster failure. To prevent the race condition from occurring, the **cman** and **qdiskd** interaction timer has been improved.

BZ#750314

Previously, a cluster partition and merge during startup fencing was not detected correctly. As a consequence, the DLM (Distributed Lock Manager) lockspace operations could become unresponsive. With this update, the partition and merge event is now detected and handled properly. DLM lockspace operations no longer become unresponsive in the described scenario.

BZ#745538

Multiple **ping** command examples on the qdisk(5) manual page did not include the **-w** option. If the **ping** command is run without the option, the action can timeout. With this update, the **-w** option has been added to those **ping** commands.

BZ#745161

Due to a bug in libgfs2, sentinel directory entries were counted as if they were real entries. As a consequence, the **mkfs.gfs2** utility created file systems which did not pass the fsck check when a large number of journal metadata blocks were required (for example, a file system with block size of 512, and 9 or more journals). With this update, incrementing the count of the directory entry is now avoided when dealing with sentinel entries. **GFS2** file systems created with large numbers of journal metadata blocks now pass the fsck check cleanly.

BZ#806002

When a node fails and gets fenced, the node is usually rebooted and joins the cluster with a fresh state. However, if a block occurs during the rejoin operation, the node cannot rejoin the cluster and the attempt fails during boot. Previously, in such a case, the cman init script did not revert actions that had happened during startup and some daemons could be erroneously left running on a node. The underlying source code has been modified so that the cman init script now performs a full rollback when errors are encountered. No daemons are left running unnecessarily in this scenario.

BZ#804938

The RELAX NG schema used to validate the cluster.conf file previously did not recognize the **totem.miss_count_const** constant as a valid option. As a consequence, users were not able to validate **cluster.conf** when this option was in use. This option is now recognized correctly by the RELAX NG schema, and the **cluster.conf** file can be validated as expected.

BZ#819787

The **cmannotifyd** daemon is often started after the **cman** utility, which means that **cmannotifyd** does not receive or dispatch any notifications on the current cluster status at startup. This update modifies the **cman** connection loop to generate a notification that the configuration and membership have changed.

BZ#749864

Incorrect use of the **free()** function in the **gfs2_edit** code could lead to memory leaks and so cause various problems. For example, when the user executed the **gfs2_edit savemeta** command, the **gfs2_edit** utility could become unresponsive or even terminate unexpectedly. This update applies multiple upstream patches so that the **free()** function is now used correctly and memory leaks no longer occur. With this update, save statistics for the **gfs2_edit savemeta** command are now reported more often so that users know that the process is still running when saving a large dinode with a huge amount of metadata.

BZ#742595

Previously, the **gfs2_grow** utility failed to expand a GFS file system if the file system contained only one resource group. This was due to the old code being based on **GFS1** (which had different fields) that calculated distances between resource groups and did not work with only one resource group. This update adds the **rgrp_size()** function in libgfs2, which calculates the size of the resource group instead of determining its distance from the previous resource group. A file system with only one resource group can now be expanded successfully.

BZ#742293

Previously, the **gfs2_edit** utility printed unclear error messages when the underlying device did not contain a valid GFS2 file system, which could be confusing. With this update, users are provided with additional information in the aforementioned scenario.

BZ#769400

Previously, the **mkfs** utility provided users with insufficient error messages when creating a **GFS2** file system. The messages also contained absolute build paths and source code references, which was unwanted. A patch has been applied to provide users with comprehensive error messages in the described scenario.

BZ#753300

The **gfs_controld** daemon ignored an error returned by the **dlm_controld** daemon for the **dlmc_fs_register()** function while mounting a file system. This resulted in a successful mount, but recovery of a **GFS** file system could not be coordinated using Distributed Lock Manager (DLM). With this update, mounting a file system is not successful under these circumstances and an error message is returned instead.

Enhancements

BZ#675723, BZ#803510

The **gfs2_convert** utility can be used on a **GFS1** file system to convert a file system from **GFS1** to **GFS2**. However, the **gfs2_convert** utility required the user to run the **gfs_fsck** utility prior to conversion, but because this tool is not included in Red Hat Enterprise Linux 6, users had to use Red Hat Enterprise Linux 5 to run this utility. With this update, the **gfs2_fsck** utility now allows users to perform a complete **GFS1** to **GFS2** conversion on Red Hat Enterprise Linux 6 systems.

BZ#678372

Cluster tuning using the **qdiskd** daemon and the **device-mapper-multipath** utility is a very complex operation, and it was previously easy to misconfigure **qdiskd** in this setup, which could consequently lead to a cluster nodes failure. Input and output operations of the **qdiskd** daemon have been improved to automatically detect multipath-related timeouts without requiring manual configuration. Users can now easily deploy **qdiskd** with device-mapper-multipath.

BZ#733298**, BZ#**740552

Previously, the **cman** utility was not able to configure Redundant Ring Protocol (RRP) correctly in corosync, resulting in RRP deployments not working propely. With this update, **cman** has been improved to configure RRP properly and to perform extra sanity checks on user configurations. It is now easier to deploy a cluster with RRP and the user is provided with more extensive error reports.

BZ#745150

With this update, Red Hat Enterprise Linux High Availability has been validated against the VMware vSphere 5.0 release.

BZ#749228

With this update, the **fence_scsi** fencing agent has been validated for use in a two-node cluster with High Availability LVM (HA-LVM).

All users of cluster and gfs2-utils are advised to upgrade to these updated package, which fix these bugs and add these enhancements.

5.27. CLUSTER-GLUE

5.27.1. RHBA-2012:0942 – cluster-glue bug fix update

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

The cluster-glue packages contain a collection of common tools that are useful for writing cluster managers such as Pacemaker.

Bug Fixes

BZ#758127

Previously, the environment variable "LRMD_MAX_CHILDREN" from the program /etc/sysconfig/pacemaker was not properly evaluated. As a result, the "max_child_count" variable in the Local Resource Management Daemon (lrmd) was not modified. With this update, the bug has been fixed so that the environment variable "LRMD_MAX_CHILDREN" is evaluated as expected.

BZ#786746

Previously, if Pacemaker attempted to cancel a recurring operation while the operation was executed, the Local Resource Management Daemon (lrmd) did not cancel the operation correctly. As a result the operation was not removed from the repeat list. With this update, a canceled operation is now marked to be removed from the repeat operation list if it is canceled during the execution so that recurring canceled operations are never executed again.

All cluster-glue users are advised to upgrade to these updated packages, which fix these bugs.

5.28. CLUSTERMON

5.28.1. RHBA-2012:0750 – 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#742431

Prior to this update, under certain circumstances, outgoing queues in inter-node communication of the modclusterd service could grow over time. To prevent this behavior, the inter-node communication is now better balanced and queues are restricted in size. Forced queue interventions are logged in the /var/log/clumond.log file.

BZ#794907

When the clustermon utility was used to get the cluster schema from the server, the schema was returned in an invalid format, preventing further processing. This bug has been fixed and clustermon now provides an exact copy of the schema in the described scenario.

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

5.29. CLUSTER

5.29.1. RHBA-2012:1480 – cluster and gfs2-utils bug fix update

Updated cluster and gfs2-utils packages that fix one bug 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 Fix

BZ#878373

Previously, the fenced daemon was creating its log file with insecure permissions. Even though no sensitive data, such as passwords, usernames, or IP addresses were ever stored in the file, with this update, log files are created with correct permissions. Permissions of an existing log file is also automatically corrected if necessary.

All users of cluster and gfs2-utils are advised to upgrade to these updated packages, which fix this bug.

5.29.2. RHBA-2012:1188 – cluster and gfs2-utils bug fix update

Updated cluster and gfs2-utils packages that fix one bug 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 Fix

BZ#849049

Previously, it was not possible to specify start-up options to the dlm_controld daemon. As a consequence, certain features were not working as expected. With this update, it is possible to use the /etc/sysconfig/cman configuration file to specify dlm_controld start-up options, thus fixing this bug.

All users of cluster and gfs2-utils are advised to upgrade to these updated packages, which fix this bug.

5.29.3. RHBA-2013:1056 – cluster and gfs2-utils bug fix update

Updated cluster and gfs2-utils packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

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 Fix

BZ#982699

Previously, the cman init script did not handle its lock file correctly. During a node reboot, this could have caused the node itself to be evicted from the cluster by other members. With this update, the cman init script now handles the lock file correctly, and no fencing action is taken by other nodes of the cluster.

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

5.30. CONMAN

5.30.1. RHEA-2012:0401 – conman enhancement update

An updated conman package that adds one enhancement is 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.

Enhancement

BZ#738967

Users are now able to configure the maximum number of open files. This allows the conman daemon to easily manage a large number of nodes.

All users of conman are advised to upgrade to this updated package, which adds this enhancement.

5.31. CONTROL-CENTER

5.31.1. RHBA-2012:0950 – control-center bug fix and enhancement update

Updated control-center packages that fix one bug and add various enhancements are now available for Red Hat Enterprise Linux 6.

The control-center packages contain various configuration utilities for the GNOME desktop. These utilities allow the user to configure accessibility options, desktop fonts, keyboard and mouse properties, sound setup, desktop theme and background, user interface properties, screen resolution, and other settings.

Bug Fix

BZ#771600

Previous versions of the control-center package contained gnome-at-mobility, a script that requires a software component that is not distributed with Red Hat Enterprise Linux 6 nor is present in any of the available channels. With this update, the non-functional gnome-at-mobility script has been removed and is no longer distributed as part of the control-center package.

Enhancements

BZ#524942

The background configuration tool now uses the XDG Base Directory Specification to determine where to store its data file. By default, this file is located at ~/.config/gnome-control-center/backgrounds.xml. Users can change the ~/.config/ prefix by setting the XDG_DATA_HOME environment variable, or set the GNOMECC_USE_OLD_BG_PATH environment variable to 1 to restore the old behavior and use the ~/.gnome2/backgrounds.xml file.

BZ#632680

The control-center-extra package now includes a GNOME Control Center shell. This shell provides a user interface for launching the various Control Center utilities.

BZ#769465, BZ#801363

The GNOME Control Center now provides a configuration utility for Wacom graphics tablets, which replaces the wacompl utility.

All users of control-center are advised to upgrade to these updated packages, which fix this bug and add these enhancements.

5.32. COOLKEY

5.32.1. RHBA-2012:0948 – coolkey bug fix update

Updated coolkey packages that resolve two issues are now available for Red Hat Enterprise Linux 6.

Coolkey is a smart card support library for the CoolKey, CAC, and PIV smart cards.

Bug Fixes

BZ#700907

Prior to this update, Coolkey did not recognize Spice virtualized CAC cards unless the card contained at least 3 certificates. This update fixes this issue so that cards with one or two certificates are recognized by Coolkey as expected. Note that this issue may also have affected some non-virtualized CAC cards.

BZ#713132

Under certain error conditions, Coolkey could leak memory data because a variable buffer was not being freed properly. With this update, the aforementioned buffer is properly freed, and memory leaks no longer occur.

All users of coolkey are advised to upgrade to these updated packages, which resolve these issues.

5.33. COREUTILS

5.33.1. RHBA-2012:0933 – coreutils bug fix and enhancement update

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

The coreutils packages contain the core GNU utilities. These packages combine the old GNU fileutils, sh-utils, and textutils packages.

Bug Fixes

BZ#772172

The "pr -c [filename]" and "pr -v [filename]" commands, which serve to show control and non-printing characters, cause the pr utility to terminate with a segmentation fault in multibyte locales. With this update, the underlying code has been modified and the pr utility now works as expected.

BZ#771843

The "-Z" option of the ls command did not explain sufficiently that only the last format option is taken into consideration and the user did not understand why the "ls -Zl" and "ls -lZ" command returned a different output. With this update, the ls info documentation has been improved.

BZ#769874

The "tail --follow" command uses the inotify API to follow the changes in a file. However, inotify does not work on remote file systems and the tail utility should fall back to polling for files on such file systems. The remote file systems GPFS and FhGFS were missing from the remote file system list and therefore "tail --follow" did not display the updates to the file on these file systems. These file systems have been added to the remote file system list and the problem no longer occurs.

BZ#751974

If SELinux was enabled, the "ls -l" command leaked one string for each non-empty directory name specified on the command line. With this update, such strings are freed from the memory and the problem no longer occurs.

BZ#754057

The su utility could remain unresponsive if it ran a process that ignored the SIG_CHLD signal. This happened because the su utility uses the waitpid() function to wait for a child process. The loop mechanism with the waitpid() function waited for the process to be in the stopped status. However, a

process masking the SIG_CHLD signal will never be in that status. With this update, the loop mechanism was improved to handle this situation correctly and the problem no longer occurs.

BZ#804604

In a non-interactive tcsh shell, the colorls.csh script returned the following error: tput: No value for $TERM and no -T specified

This happened because the tcsh shell did not short-circuit the evaluation of the logical AND in a colorls.csh expression. With this update, checking for an interactive shell has been modified and the script no longer returns the error message.

Enhancements

BZ#766461

In the default listing, the df utility showed long file system names including UUID. Consequently, the columns following the file system names were pushed to the right and made the df output hard to read. As long UUID system names are becoming more common, df now prints the referent when a long name refers to a symlink, and no file systems are specified.

BZ#691466

The user could not use octal digit mode when cleaning special set-user-id and set-group-id bits on a directory with the chmod tool. This is an upstream change, however as it was possible in all the previous Red Hat Enterprise Linux releases, it is necessary to provide backwards compatibility. Therefore, the chmod tool now again allows the user to clear the special bits on the directories using octal digit mode if the octal digit mode is at least 5 digits long.

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

5.34. COROSYNC

5.34.1. RHBA-2012:1237 – corosync bug fix update

Updated corosync packages that fix a bug 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 Fix

BZ#849554

Previously, the corosync-notifyd daemon, with dbus output enabled, waited 0.5 seconds each time a message was sent through dbus. Consequently, corosync-notifyd was extremely slow in producing output and memory of the Corosync server grew. In addition, when corosync-notifyd was killed, its memory was not freed. With this update, corosync-notifyd no longer slows down its operation with these half-second delays and Corosync now properly frees memory when an IPC client exits.

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

5.34.2. RHBA-2012:0777 – corosync bug fix and enhancement update

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

The corosync packages provide the Corosync Cluster Engine and the C language APIs for Red Hat Enterprise Linux cluster software.

Bug Fixes

BZ#741455

The mainconfig module passed an incorrect string pointer to the function that opens the corosync log file. If the path to the file (in cluster.conf) contained a non-existing directory, an incorrect error message was returned stating that there was a configuration file error. The correct error message is now returned informing the user that the log file cannot be created.

BZ#797192

The coroipcc library did not delete temporary buffers used for Inter-Process Communication (IPC) connections that are stored in the /dev/shm shared-memory file system. The /dev/shm memory resources became fully used and caused a Denial of Service event. The library has been modified so that applications delete temporary buffers if the buffers were not deleted by the corosync server. The /dev/shm system is now no longer cluttered with needless data.

BZ#758209

The range condition for the update_aru() function could cause incorrect checking of message IDs. The corosync utility entered the "FAILED TO RECEIVE" state and failed to receive multicast packets. The range value in the update_aru() function is no longer checked and the check is now performed using the fail_to_recv_const constant.

BZ#752159

If the corosync-notifyd daemon was running for a long time, the corosync process consumed an excessive amount of memory. This happened because the corosync-notifyd daemon failed to indicate that the no-longer used corosync objects were removed, resulting in memory leaks. The corosync-notifyd daemon has been fixed and the corosync memory usage no longer increases if corosync-notifyd is running for long periods of time.

BZ#743813

When a large cluster was booted or multiple corosync instances started at the same time, the CPG (Closed Process Group) events were not sent to the user. Therefore, nodes were incorrectly detected as no longer available, or as leaving and re-joining the cluster. The CPG service now checks the exit code in such scenarios properly and the CPG events are sent to users as expected.

BZ#743815

The OpenAIS EVT (Eventing) service sometimes caused deadlocks in corosync between the timer and serialize locks. The order of locking has been modified and the bug has been fixed.

BZ#743812

When corosync became overloaded, IPC messages could be lost without any notification. This happened because some services did not handle the error code returned by the totem_mcast() function. Applications that use IPC now handle the inability to send IPC messages properly and try sending the messages again.

BZ#747628

If both the corosync and cman RPM packages were installed on one system, the RPM verification process failed. This happened because both packages own the same directory but apply different rights to it. Now, the RPM packages have the same rights and the RPM verification no longer fails.

BZ#752951

corosync consumed excessive memory because the getaddrinfo() function leaked memory. The memory is now freed using the freeadrrinfo() function and getaddrinfo() no longer leaks memory.

BZ#773720

It was not possible to activate or deactivate debug logs at runtime due to memory corruption in the objdb structure. The debug logging can now be activated or deactivated on runtime, for example by the "corosync-objctl -w logging.debug=off" command.

Enhancement

BZ#743810

Each IPC connection uses 48 K in the stack. Previously, multi-threading applications with reduced stack size did not work correctly, which resulted in excessive memory usage. Temporary memory resources in a heap are now allocated to the IPC connections so that multi-threading applications no longer need to justify IPC connections' stack size.

All users of corosync are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.34.3. RHBA-2013:0731 – corosync bug fix update

Updated corosync packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

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

Bug Fix

BZ#929100

When running applications which used the Corosync IPC library, some messages in the dispatch() function were lost or duplicated. This update properly checks the 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. Now, messages in the dispatch() function are no longer lost or duplicated.

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

5.35. CPIO

5.35.1. RHBA-2012:1414 – cpio bug fix update

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

The cpio packages provide the GNU cpio utility for creating and extracting archives, or copying files from one place to another.

Bug Fix

BZ#866467

Previously, the cpio command was unable to split file names longer than 155 bytes into two parts during the archiving operation. Consequently, cpio could terminate unexpectedly with a segmentation fault. This bug has been fixed and cpio now handles long file names without any crashes.

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

5.35.2. RHBA-2012:0444 – cpio bug fix update

An updated cpio package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The cpio package provides the GNU cpio file archiver utility. GNU cpio can be used to copy and extract files into or from cpio and Tar archives.

Bug Fix

BZ#746209

Prior to this update,the options --to-stdout and --no-absolute-filenames were not listed in the cpio (1) manual page. This update includes the missing options and corrects several misprints.

All users of cpio are advised to upgrade to this updated package, which fixes this bug.

5.36. CPUSPEED

5.36.1. RHBA-2012:1404 – cpuspeed bug fix update

Updated cpuspeed packages that fix four bugs are now available for Red Hat Enterprise Linux 6.

The cpuspeed packages provide a daemon to manage the CPU frequency scaling.

Bug Fixes

BZ#642838

Prior to this update, the PCC driver used the “userspace” governor was loaded instead of the “ondemand” governor when loading. This update modifies the init script to also check the PCC driver.

BZ#738463

Prior to this update, the cpuspeed init script tried to set cpufrequency system files on a per core basis which was a deprecated procedure. This update sets thresholds globally.

BZ#616976

Prior to this update, the cpuspeed tool did not reset MIN and MAX values, when the configuration file was emptied. As a consequence, the MIN_SPEED or MAX_SPEED values were not reset as expected. This update adds conditionals in the init script to check these values. Now, the MIN_SPEED or MAX_SPEED values are reset as expected.

BZ#797055

Prior to this update, the init script did not handle the IGNORE_NICE parameter as expected. As a consequence, "-n" was added to command options when the IGNORE_NICE parameter was set. This update modifies the init script to stop adding the NICE option when using the IGNORE_NICE parameter.

All users of cpuspeed are advised to upgrade to these updated packages, which fix these bugs.

5.37. CRASH

5.37.1. RHBA-2012:0822 – crash bug fix and enhancement update

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

The crash package provides 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.

The crash package has been upgraded to upstream version 6.0.4, which provides a number of bug fixes and enhancements over the previous version. (BZ#767257)

Bug Fixes

BZ#754291

If the kernel was configured with the Completely Fair Scheduler (CFS) Group Scheduling feature enabled (CONFIG_FAIR_GROUP_SCHED=y), the "runq" command of the crash utility did not display all tasks in CPU run queues. This update modifies the crash utility so that all tasks in run queues are now displayed as expected. Also, the "-d" option has been added to the "runq" command, which provides debugging information same as the /proc/sched_debug file.

BZ#768189

The "bt" command previously did not handle recursive non-maskable interrupts (NMIs) correctly on the Intel 64 and AMD64 architectures. As a consequence, the "bt" command could, under certain circumstances, display a task backtrace in an infinite loop. With this update, the crash utility has been modified to recognize a recursion in the NMI handler and prevent the infinite displaying of a backtrace.

BZ#782837

Under certain circumstances, the number of the "elf_prstatus" entries in the header of the compressed kdump core file could differ from the number of CPUs running when the system crashed. If such a core file was analyzed by the crash utility, crash terminated unexpectedly with a segmentation fault while displaying task backtraces. This update modifies the code so that the "bt" command now displays a backtrace as expected in this scenario.

BZ#797229

Recent changes in the code caused the crash utility to incorrectly recognize compressed kdump dump files for the 64-bit PowerPC architecture as dump files for the 32-bit PowerPC architecture. This caused the crash utility to fail during initialization. This update fixes the problem and the crash utility now recognizes and analyzes the compressed kdump dump files for the 32-bit and 64-bit PowerPC architectures as expected.

BZ#817247

The crash utility did not correctly handle situations when a user page was either swapped out or was not mapped on the IBM System z architecture. As a consequence, the "vm -p" command failed and either a read error occurred or an offset va1lue of a swap device was set incorrectly. With this update, crash displays the correct offset value of the swap device or correctly indicates that the user page is not mapped.

BZ#817248

The crash utility did not correctly handle situations when the "bt -t" and "bt -T" commands were run on an active task on a live system on the IBM System z architecture. Consequently, the commands failed with the "bt: invalid/stale stack pointer for this task: 0" error message. This update modifies the source code so that the "bt -t" and "bt -T" commands execute as expected.

Enhancements

BZ#736884

With this update, crash now supports the "sadump" dump file format created by the Fujitsu Stand Alone Dump facility.

BZ#738865

The crash utility has been modified to fully support the "ELF kdump" and "compressed kdump" dump file formats for IBM System z.

BZ#739096

The makedumpfile facility can be used to filter out specific kernel data when creating a dump file, which can cause the crash utility to behave unpredictably. With this update, the crash utility now displays an early warning message if any part of the kernel has been erased or filtered out by makedumpfile.

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

5.38. CRASH-TRACE-COMMAND

5.38.1. RHBA-2012:0808 – crash-trace-command bug fix update

An updated crash-trace-command package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The crash-trace-command package provides a trace extension module for the crash utility, allowing it to read ftrace data from a core dump file.

Bug Fix

BZ#729018

Previously, the "trace.so" binary in the crash-trace-command package was compiled by the GCC compiler without the "-g" option. Therefore, no debugging information was included in its associated "trace.so.debug" file. This could affect a crash analysis performed by the Automatic Bug Reporting Tool (ABRT) and its retrace server. Also, proper debugging of crashes using the GDB utility was not possible under these circumstances. This update modifies the Makefile of crash-trace-command to

compile the "trace.so" binary with the "RPM_OPT_FLAGS" flag, which ensures that the GCC's "-g" option is used during the compilation. Debugging and a crash analysis can now be performed as expected.

All users of crash-trace-command are advised to upgrade to this updated package, which fixes this bug.

5.39. CREATEREPO

5.39.1. RHBA-2012:0354 – createrepo bug fix and enhancement update

An updated createrepo package that fixes one bug is now available for Red Hat Enterprise Linux 6.

This package contains scripts that generate a common metadata repository from a directory of RPM packages.

Bug Fix

BZ#623105

Prior to this update, the shebang line of the modifyrepo.py script contained "#!/usr/bin/env python", so the system path was used to locate the Python executable. When another version of Python was installed on the system, and "/usr/local/python" was specified in the PATH environment variable, scripts did not work due to Python compatibility problems. With this update, the shebang line is modified to "#!/usr/bin/python", so that the system version of Python is always used.

All users of createrepo are advised to upgrade to this updated package, which fixes this bug.

5.40. CRYPTSETUP-LUKS

5.40.1. RHBA-2012:0886 – cryptsetup-luks bug fix update

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

The cryptsetup-luks packages provide a utility which allows users to set up encrypted devices with the Device Mapper and the dm-crypt target.

Bug Fix

BZ#746648

For some configurations, the cryptsetup utility incorrectly translated major:minor device pairs to device names in the /dev/ directory (for example, on HP Smart Array devices). With this update, the underlying source code has been modified to address this issue, and the cryptsetup utility now works as expected. (BZ#755478) * If a device argument for the "cryptsetup status" command included a /dev/mapper/ prefix, the prefix was duplicated in the command's output. The output was fixed and no longer includes duplicated strings.

All users of cryptsetup-luks are advised to upgrade to these updated packages, which fix these bugs.

5.41. CTDB

5.41.1. RHBA-2012:0904 – ctdb bug fix update

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

The ctdb packages provide a clustered database based on Samba's Trivial Database (TDB) used to store temporary data.

Bug Fix

BZ#794888

Prior to this update, the ctdb working directory, all subdirectories and the files within were created with incorrect SELinux contexts when the ctdb service was started. This update uses the post-install script to create the ctdb directory, and the command "/sbin/restorecon -R /var/ctdb" sets now the right SELinux context.

All users of ctdb are advised to upgrade to these updated packages, which fix this bug.

5.42. CUPS

5.42.1. RHBA-2012:1285 – cups bug fix update

Updated cups packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar operating systems.

Bug Fix

BZ#854472

Previously, when no authentication was initially provided (or even requested), cups returned the "forbidden" status rather than the correct "unauthorized" status. Consequently, certain operations, such as attempts to move a job between queues using the web user interface, failed. An upstream patch has been provided to address this bug and cups now returns correct status in the described scenario.

All users of cups are advised to upgrade to these updated packages, which fix this bug.

5.42.2. RHBA-2012:1470 – cups bug fix update

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

The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar operating systems.

Bug Fix

BZ#873592

Previously, with LDAP browsing enabled, one of the objects used for LDAP queries was freed twice, which caused the cupsd service to terminate unexpectedly with a segmentation fault. Additionally, names of browsed LDAP queues were truncated by a single character. Consequently, only one print queue was listed if multiple print queues with names varying only in the last character were defined. With this update, an upstream patch that resolves these problems has been back-ported, and the cupsd service no longer crashes and LDAP print queues are now displayed correctly.

All users of cups are advised to upgrade to these updated packages, which fix this bug.

5.42.3. RHBA-2012:0818 – cups bug fix update

Updated cups packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar operating systems.

Bug Fixes

BZ#738410

Prior to this update, the textonly filter did not always correctly generate output when a single copy was requested. The textonly filter generates output for a single or multiple copies by spooling the output for one copy into a temporary file, then sending the content of that temporary file as many times as required. However, if the filter was used for the MIME-type conversion rather than as a PostScript Printer Description (PPD) filter, and a single copy was requested, the temporary file was not created and the program failed with the "No such file or directory" message. With this update, the textonly filter has been modified to create a temporary file regardless of the number of copies specified. The data is now sent to the printer as expected.

BZ#738914, BZ#740093

Previously, empty jobs could be created using the "lp" command either by submitting an empty file to print (for example by executing "lp /dev/null") or by providing an empty file as standard input. In this way, a job was created but was never processed. With this update, creation of empty print jobs is not allowed, and the user is now informed that no file is in the request.

BZ#806818

The German translation for the search page template of the web interface contained an error that prevented the search feature from functioning correctly: attempting to search for a printer in the CUPS web interface failed, and an error message was displayed in the browser. The bug in the search template has been fixed, and the search feature in the German locale now works as expected in this scenario.

All users of cups are advised to upgrade to these updated packages, which fix these bugs.

5.43. CVS

5.43.1. RHBA-2012:1302 – cvs bug fix update

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

[Update 19 November 2012] The file list of this advisory was updated to move the new cvs-inetd package from the base repository to the optional repository in the Client and HPC Node variants. No changes have been made to the packages themselves.

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.

* Prior to this update, the C shell (csh) did not set the CVS_RSH environment variable to "ssh" and the remote shell (rsh) was used instead when the users accessed a remote CVS server. As a consequence, the connection was vulnerable to attacks because the remote shell is not encrypted or not necessarily

enabled on every remote server. The cvs.csh script now uses valid csh syntax and the CVS_RSH environment variable is properly set at log-in. (BZ#671145)

* Prior to this update, the xinetd package was not a dependency of the cvs package. As a result, the CVS server was not accessible through network. With this update, the cvs-inetd package, which contains the CVS inetd configuration file, ensures that the xinetd package is installed as a dependency and the xinetd daemon is available on the system. (BZ#695719)

Bug Fixes

BZ#671145

Prior to this update, the C shell (csh) did not set the CVS_RSH environment variable to "ssh" and the remote shell (rsh) was used instead when the users accessed a remote CVS server. As a consequence, the connection was vulnerable to attacks because the remote shell is not encrypted or not necessarily enabled on every remote server. The cvs.csh script now uses valid csh syntax and the CVS_RSH environment variable is properly set at log-in.

BZ#695719

Prior to this update, the xinetd package was not a dependency of the cvs package. As a result, the CVS server was not accessible through network. With this update, the cvs-inetd package, which contains the CVS inetd configuration file, ensures that the xinetd package is installed as a dependency and the xinetd daemon is available on the system.

All users of cvs are advised to upgrade to these updated packages, which fix these bugs.

5.44. CYRUS-SASL

5.44.1. RHBA-2012:1495 – cyrus-sasl bug fix update

Updated cyrus-sasl packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The cyrus-sasl packages contain the Cyrus implementation of Simple Authentication and Security Layer (SASL). SASL is a method for adding authentication support to connection-based protocols.

Bug Fix

BZ#878357

Previously, the GSSAPI plug-in kept credential handles open the whole time a client was connected. These handles hold a pointer to a Kerberos replay cache structure. When the replay cache is a file, that structure includes an open file descriptor. When too many clients were using GSSAPI, the server could run out of file handles. Consequently, the client could become unresponsive until restarted. With this update, a GSSAPI credential handle is closed immediately after the plug-in gets the security context, thus preventing this bug.

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

5.45. DASH

5.45.1. RHBA-2012:1381 – dash bug fix update

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

The dash packages provide the POSIX-compliant Debian Almquist shell intended for small media like floppy disks.

Bug Fix

BZ#706147

Prior to this update, the dash shell was not an allowed login shell. As a consequence, users could not log in using the dash shell. This update adds the dash to the /etc/shells list of allowed login shells when installing or upgrading dash package and removes it from the list when uninstalling the package. Now, users can login using the dash shell.

All users of dash are advised to upgrade to these updated packages, which fix this bug.

5.46. DB4

5.46.1. RHBA-2012:0452 – db4 bug fix update

An updated db4 package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The Berkeley Database (Berkeley DB) is a programmatic toolkit that provides embedded database support for both traditional and client/server applications. The Berkeley DB includes B+tree, Extended Linear Hashing, Fixed and Variable-length record access methods, transactions, locking, logging, shared memory caching, and database recovery. The Berkeley DB supports C, C++, Java, and Perl APIs. It is used by many applications, including Python and Perl, so this should be installed on all systems.

Bug Fix

BZ#784662

The db4 spec file incorrectly stated that the "License" is simply "BSD", whereas it is in fact licensed under both the BSD and Sleepycat licenses, the latter of which differs from the Berkeley Software Distribution (BSD) license by including a redistribution clause. This update corrects the spec file so it correctly states that the db4 software is provided under the "Sleepycat and BSD" license.

Users of db4 are advised to upgrade to this updated package which fixes this bug.

5.46.2. RHBA-2013:1444 – db4 bug fix update

Updated db4 packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The Berkeley Database (Berkeley DB) is a programmatic toolkit that provides embedded database support for both traditional and client/server applications. The Berkeley DB includes B+tree, Extended Linear Hashing, Fixed and Variable-length record access methods, transactions, locking, logging, shared memory caching, and database recovery. The Berkeley DB supports C, C++, Java, and Perl APIs. It is used by many applications, including Python and Perl, so this should be installed on all systems.

Bug Fix

BZ#1012586

Due to an incorrect order of the mutex initialization calls, the rpm utility became unresponsive under certain circumstances, until it was terminated. With this update, the order of the mutex initialization calls has been revised. As a result, the rpm utility no longer becomes unresponsive.

Users of db4 are advised to upgrade to these updated packages, which fix this bug.

5.47. DBUS

5.47.1. RHSA-2012:1261 – Moderate: dbus security update

Updated dbus packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

D-Bus is a system for sending messages between applications. It is used for the system-wide message bus service and as a per-user-login-session messaging facility.

Security Fix

CVE-2012-3524

It was discovered that the D-Bus library honored environment settings even when running with elevated privileges. A local attacker could possibly use this flaw to escalate their privileges, by setting specific environment variables before running a setuid or setgid application linked against the D-Bus library (libdbus).

Note: With this update, libdbus ignores environment variables when used by setuid or setgid applications. The environment is not ignored when an application gains privileges via file system capabilities; however, no application shipped in Red Hat Enterprise Linux 6 gains privileges via file system capabilities.

Red Hat would like to thank Sebastian Krahmer of the SUSE Security Team for reporting this issue.

All users are advised to upgrade to these updated packages, which contain a backported patch to correct this issue. For the update to take effect, all running instances of dbus-daemon and all running applications using the libdbus library must be restarted, or the system rebooted.

5.48. DEVICE-MAPPER-MULTIPATH

5.48.1. RHBA-2012:1111 – device-mapper-multipath bug fix update

Updated device-mapper-multipath packages that fix one bug 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 Fix

BZ#837594

When a multipath vector (a dynamically allocated array) was resized to a smaller size, device-

mapper-multipath did not reassign the pointer to the array. If the array location was changed by reducing its size, device-mapper-multipath corrupted its memory. With this update, device-mapper-multipath correctly reassigns the pointer in this scenario, and memory corruption no longer occurs.

All users of device-mapper-multipath are advised to upgrade to these updated packages, which fix this bug.

5.48.2. RHBA-2012:0946 – 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 to manage multipath devices using the device-mapper multipath kernel module.

Bug Fixes

BZ#812832

The **multipathd** daemon was not correctly stopping waiter threads during shutdown. The waiter threads could access freed memory and cause the daemon to terminate unexpectedly during shutdown. With this update, the **mutlipathd** daemon now correctly stops the waiter threads before they can access any freed memory and no longer crashes during shutdown.

BZ#662433

When **Device Mapper Multipath** was stopped, **multipathd** did not disable the **queue_if_no_path** option on multipath devices by default. When **multipathd** was stopped during shutdown, I/O of the device was added to the queue if all paths to a device were lost, and the shutdown process became unresponsive. With this update, multipathd now sets the **queue_without_daemon** option to **no** by default. As a result, all multipath devices stop queueing when **multipathd** is stopped and multipath now shuts down as expected.

BZ#752989

Device Mapper Multipath uses regular expressions in built-in device configurations to determine a multipath device so as to apply the correct configuration to the device. Previously, some regular expressions for resolving the device vendor name and product ID were not specific enough. As a consequence, some devices could be matched with incorrect device configurations. With this update, the product and vendor regular expressions have been modified so that all multipath devices are now configured properly.

BZ#754586

After renaming a device, there was a race condition between **multipathd** and udev to rename the new multipath device nodes. If udev renamed the device node first, **multipathd** removed the device created by udev and failed to create the new device node. With this update, **multipathd** immediately creates the new device nodes, and the race condition no longer occurs. As a result, the renamed device is now available as expected.

BZ#769527

Previously, the **flush_on_last_dev** handling code did not implement handling of the queue feature properly. Consequently, even though the **flush_on_last_del** feature was activated, **multipathd** re-enabled queueing on multipath devices that could not be removed immediately after the last path

device was deleted. With this update, the code has been fixed and when the user sets **flush_on_last_del**, their multipath devices correctly disable queueing, even if the devices cannot be closed immediately.

BZ#796384

Previously, **Device Mapper Multipath** used a fixed-size buffer to read the Virtual Device Identification page [0x83]. The buffer size was sometimes insufficient to accommodate the data sent by devices and the ALUA (Asymmetric Logical Unit Access) prioritizer failed. Device Mapper Multipath now dynamically allocates a buffer large enough for the Virtual Device Identification page and the ALUA prioritizer no longer fails in the scenario described.

BZ#744210

Previously, **multipathd** did not set the **max_fds** option by default, which sets the maximum number of file descriptors that **multipathd** can open. Also, the **user_friendly_names** setting could only be configured in the **defaults** section of **/etc/multipath.conf**. The user had to set **max_fds** manually and override the **default user_friendly_names** value in their device-specific configurations. With this update, multipath now sets **max_fds** to the system maximum by default, and **user_friendly_names** can be configured in the **devices** section of **multipath.conf**. Users no longer need to set max_fds for large setups, and they are able to select user_friendly_names per device type.

BZ#744756

Previously, to modify a built-in configuration, the vendor and product strings of the user's configuration had to be identical to the vendor and product strings of the built-in configuration. The vendor and product strings are regular expressions, and the user did not always know the correct vendor and product strings needed to modify a built-in configuration. With this update, the **hwtable_regex_match** option was added to the defaults section of **multipath.conf**. If it is set to **yes**, Multipath uses regular-expression matching to determine if the user's vendor and product strings match the built-in device configuration strings: the user can use the actual vendor and product information from their hardware in their device configuration, and it will modify the default configuration for that device. The option is set to **no** by default.

BZ#750132

Previously, **multipathd** was using a deprecated Out-of-Memory (OOM) adjustment interface. Consequently, the daemon was not protected from the OOM killer properly; the OOM killer could kill the daemon when memory was low and the user was unable to restore failed paths. With this update, **multipathd** now uses the new Out-of-Memory adjustment interface and can no longer be killed by the Out-of-Memory killer.

BZ#702222

The **multipath.conf** file now contains a comment which informs the user that the configuration must be reloaded for any changes to take effect.

BZ#751938

The **multipathd** daemon incorrectly exited with code **1** when **multipath -h** (print usage) was run. With this update, the underlying code has been modified and **multipathd** now returns code **0** as expected in the scenario described.

BZ#751039

Some **multipathd** threads did not check if **multipathd** was shutting down before they started their execution. Consequently, the **multipathd** daemon could terminate unexpectedly with a segmentation fault on shutdown. With this update, the **multipathd** threads now check if **multipathd**

is shutting down before triggering their execution, and **multipathd** no longer terminates with a segmentation fault on shutdown.

BZ#467709

The **multipathd** daemon did not have a failover method to handle switching of path groups when multiple nodes were using the same storage. Consequently, if one node lost access to the preferred paths to a logical unit, while the preferred path of the other node was preserved, **multipathd** could end up switching back and forth between path groups. This update adds the **followover** failback method to device-mapper-multipath. If the **followover** failback method is set, **multipathd** does not fail back to the preferred path group, unless it just came back online. When multiple nodes are using the same storage, a path failing on one machine now no longer causes the path groups to continually switch back and forth.

Enhancements

BZ#737051

The **NetApp** brand name has been added to the documentation about the RDAC (Redundant Disk Array Controller) checker and prioritizer.

BZ#788963

The built-in device configuration for **Fujitsu ETERNUS** has been added.

BZ#760852

If the multipath checker configuration was set to **tur**, the checks were not performed asynchronously. If a device failed and the checker was waiting for the SCSI layer to fail back, the checks on other paths were kept waiting. The checker has been rewritten so as to check the paths asynchronously, and the path checking on other paths continues as expected.

BZ#799908

A built-in configuration for **IBM XIV** Storage System has been added.

BZ#799842

The NetApp LUN built-in configuration now uses the **tur** path checker by default. Also flush_on_last_del has been enabled, dev_loss_tmo has been set to **infinity**, fast_io_fail_tmo has been set to **5**, and pg_init_retries has been set to **50**.

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

5.49. DHCP

5.49.1. RHSA-2012:1141 – Moderate: dhcp security update

Updated dhcp packages that fix three security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

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.

Security Fixes

CVE-2012-3571

A denial of service flaw was found in the way the dhcpd daemon handled zero-length client identifiers. A remote attacker could use this flaw to send a specially-crafted request to dhcpd, possibly causing it to enter an infinite loop and consume an excessive amount of CPU time.

CVE-2012-3954

Two memory leak flaws were found in the dhcpd daemon. A remote attacker could use these flaws to cause dhcpd to exhaust all available memory by sending a large number of DHCP requests.

Upstream acknowledges Markus Hietava of the Codenomicon CROSS project as the original reporter of CVE-2012-3571, and Glen Eustace of Massey University, New Zealand, as the original reporter of CVE-2012-3954.

Users of DHCP should upgrade to these updated packages, which contain backported patches to correct these issues. After installing this update, all DHCP servers will be restarted automatically.

5.49.2. RHBA-2012:0793 – dhcp bug fix and enhancement update

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

The dhcp package provides software to support the Dynamic Host Configuration Protocol (DHCP) and DHCPv6 protocol. The Dynamic Host Configuration Protocol (DHCP) is a protocol that allows individual devices on an IP network to obtain their own network configuration information, including an IP address, a subnet mask, and a broadcast address.

Bug Fixes

BZ#656339

Previously, when dhclient was unsuccessful in obtaining or renewing an address, it restored the resolv.conf file from backup even when there were other dhclient processes running. Consequently, network traffic could be unnecessarily interrupted. The bug in dhclient-script has been fixed and dhclient now restores resolv.conf from backup only if there are no other dhclient processes running.

BZ#747017

A bug caused an infinite loop in a dhcpd process when dhcpd tried to parse the slp-service-scope option in dhcpd.conf. As a consequence, dhcpd entered an infinite loop on startup consuming 100% of the CPU cycles. This update improves the code and the problem no longer occurs.

BZ#752116

Previously, the DHCPv4 client did not check whether the address received in a DHCPACK message was already in use. As a consequence, it was possible that after a reboot two clients could have the same, conflicting, IP address. With this update, the bug has been fixed and DHCPv4 client now performs duplicate address detection (DAD) and sends a DHCPDECLINE message if the address received in DHCPACK is already in use, as per RFC 2131.

BZ#756759

When dhclient is invoked with the "-1" command-line option, it should try to get a lease once and on failure exit with status code 2. Previously, when dhclient was invoked with the "-1" command-line option, and then issued a DHCPDECLINE message, it continued in trying to obtain a lease. With this update, the dhclient code has been fixed. As a result, dhclient stops trying to obtain a lease and exits after sending DHCPDECLINE when started with the "-1" option.

BZ#789719

Previously, dhclient kept sending DHCPDISCOVER messages in an infinite loop when started with the "-timeout" option having a value of 3 or less (seconds). With this update, the problem has been fixed and the "-timeout" option works as expected with all values.

Enhancements

BZ#790686

The DHCP server daemon now uses portreserve for reserving ports 647 and 847 to prevent other programs from occupying them.

BZ#798735

All DHCPv6 options defined in RFC5970, except for the Boot File Parameters Option, were implemented. This allows the DHCPv6 server to pass boot file URLs back to IPv6-based netbooting clients (UEFI) based on the system's architecture.

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

5.50. DING-LIBS

5.50.1. RHBA-2012:0799 – ding-libs bug fix update

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

The ding-libs packages contain a set of libraries used by the System Security Services Daemon (SSSD) and other projects and provide functions to manipulate filesystem pathnames (libpath_utils), a hash table to manage storage and access time properties (libdhash), a data type to collect data in a hierarchical structure (libcollection), a dynamically growing, reference-counted array (libref_array), and a library to process configuration files in initialization format (INI) into a library collection data structure (libini_config).

Bug Fixes

BZ#736074

Prior to this update, memory could become corrupted if the initial table size exceeded 1024 buckets. This update modifies libdhash so that large initial table sizes now correctly allocate memory.

BZ#801393

Prior to this update, buffers were filled and one character above the allocated size would be set to the null terminator if the combination of two strings,concatenated by the function path_concat(), exceeded the size of the destination buffer. This update modifies the underlying code so that the null terminator is no longer added after the end of the buffer.

All users of ding-libs are advised to upgrade to these updated packages, which fix these bugs.

5.51. DMRAID

5.51.1. RHBA-2012:0910 – dmraid bug fix update

Updated dmraid packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The dmraid packages provide the ATARAID/DDF1 activation tool. The tool supports RAID device discovery and RAID set activation. It also displays properties for ATARAID/DDF1-formatted RAID sets on Linux kernels using the device-mapper utility.

Bug Fixes

BZ#729971

Prior to this update, a grub installation failed silently on a dmraid mirror because the device geometry of RAID sets was not set properly. Consequently, the set partition's MBR failed to be created and the partition failed to boot. With this update, the underlying code has been modified and the geometry on dmraid devices is set up correctly.

BZ#729032

The dmraid binary was compiled without gcc's -g option and the debuginfo file did not contain the ".debug_info" section. Consequently, it was not possible to generate debugging information and debug dmraid properly. With this update, the binary has been compiled with the proper debugging options and the problem no longer occurs.

BZ#701501

When the dmraid tool was accessing a 4 KB sector or smaller, it returned a misleading error message. With this update, the library function that checks the device size has been modified and the error message is no longer displayed under these circumstances.

All users of dmraid are advised to upgrade to these updated packages, which fix these bug.

5.52. DNSMASQ

5.52.1. RHEA-2012:0869 – dnsmasq enhancement update

Updated dnsmasq packages that add an enhancement are now available for Red Hat Enterprise Linux 6.

The dnsmasq package contains Dnsmasq, a lightweight DNS (Domain Name Server) forwarder and DHCP (Dynamic Host Configuration Protocol) server.

Enhancement

BZ#794792

A new subpackage, dnsmasq-utils, has been added. The dnsmasq-utils subpackage contains the dhcp_lease_time and dhcp_release utilities, which serve to query and remove DHCP server leases using the standard DHCP protocol.

All dnsmasq users are advised to upgrade to these updated packages, which add this enhancement.

5.53. DOCBOOK-UTILS

5.53.1. RHBA-2012:1321 – docbook-utils bug fix update

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

The docbook-utils packages provide a set of utility scripts to convert and analyze SGML documents in general, and DocBook files in particular. The scripts are used to convert from DocBook or other SGML formats into file formats like HTML, man, info, RTF and many more.

Bug Fixes

BZ#639866

Prior to this update, the Perl script used for generating manpages contained a misprint in the header. As a consequence, the header syntax of all manual pages that docbook-utils built was wrong. This update corrects the script. Now the manual page headers have the right syntax.

All users of docbook-utils are advised to upgrade to these updated packages, which fix this bug.

5.54. DRACUT

5.54.1. RHBA-2012:1318 – dracut bug fix update

Updated dracut packages that fix a bug 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 Fix

BZ#860351

If the "/boot/" directory was not on a separate file system, dracut called the sha512hmac utility with a file name prefixed with "/sysroot/boot". Consequently, sha512mac searched for the file checksum in "/boot/", returned errors, and dracut considered the FIPS check to have failed. Eventually, a kernel panic occurred. With this update, dracut uses a symlink linking "/boot" to "/sysroot/boot", sha512mac can now access files in "/boot/", and FIPS checks now pass, allowing the system to boot properly in the described scenario.

All users of dracut are advised to upgrade to these updated packages, which fix this bug.

5.54.2. RHBA-2012:1078 – dracut bug fix update

Updated dracut packages that fix one bug 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.

Bug Fix

BZ#839296

Previously, the default mount option of the proc file system used during boot was "mount -t proc -o nosuid,noexec,nodev proc /proc". This caused that device nodes in the proc file system were inaccessible by certain kernel drivers. With this update, the option has been changed to previously used "mount -t proc proc /proc", so that the proc file system can be successfully accessed by kernel drivers.

All users of dracut are advised to upgrade to these updated packages, which fix this bug.

5.54.3. RHBA-2012:0839 – dracut bug fix and enhancement update

Updated dracut packages that fix multiple bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The dracut packages include an event-driven initramfs generator infrastructure based on udev. 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#788119

Previously, if a dracut module did not contain an "install" file, dracut could not execute the "installkernel" command. Consequently, the dracut fips-aesni module could not be included in the initramfs image. Now, "installkernel" can be correctly executed in the described scenario, thus fixing this bug.

BZ#761584

Previously, dracut failed to start up a degraded RAID array, resulting in a non-booting system. With this update, dracut uses the rd_retry kernel command-line parameter value and after rd_retry/2 seconds attempts to force the array to start, thus fixing this bug.

BZ#747840

During boot-up, dracut called the "udevadm settle" command several times. As a result, inconsequential messages about the command timeout were sometimes returned, creating clutter in the console output. This update fixes the bug and the messages are no longer returned in the described scenario.

BZ#735529

Occasionally, dracut attempted to assemble an array before all disks were available. As a result, dracut started the array in degraded mode or failed altogether. This bug has been fixed and dracut now forces degraded arrays to start only after a period of time controlled by the rd_retry kernel command-line parameter.

BZ#714039

The dracut package depended on the vconfig package although vconfig is not used by dracut. This update removes the dependency on vconfig.

BZ#794863

Previously, if a network interface was brought up, dracut waited for two seconds to detect that the link was up. For certain network cards, two seconds is not long enough. Consequently, the network was not properly set up and the system could not boot. Now, dracut waits for ten seconds, thus fixing this bug.

BZ#752584

Dracut did not set the broadcast address for network interfaces it started up, resulting in a 0.0.0.0 broadcast address. This bug has been fixed and the default broadcast address is now set properly on startup.

BZ#703164

The FILES section of the dracut man page has been amended to fix inaccurate content.

BZ#752073

If the user adds multiple "console=[tty]" parameters on the kernel command line, the last parameter specifies the primary console. Previously, dracut failed to initialize this console and instead initialized /dev/tty0 unconditionally. This bug has been fixed and dracut now initializes the correct console in the described scenario.

BZ#788618

When no user name and password were specified in an iSCSI interface, dracut reused the login information from a previous iSCSI parameter. Consequently, the authentication failed and the system did not boot up. This update fixes the bug.

Enhancements

BZ#722879

Previously, it was not possible to exclude a kernel driver from the initramfs image to reduce its size. This update introduces the "--omit-driver" option to provide this functionality.

BZ#752005

The "lsinitrd" command has been enhanced to support initramfs images compressed by the LZMA algorithm.

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

5.55. DROPWATCH

5.55.1. RHBA-2012:1182 – dropwatch bug fix update

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

The dropwatch package contains a utility that provides packet monitoring services.

Bug Fix

BZ#725464

Prior to this update, the dropwatch utility could become unresponsive because it was waiting for a deactivation acknowledgement to be issued by an already deactivated or stopped service. With this update, dropwatch detects an attempt to deactivate/stop an already deactivated/stopped service and no longer hangs.

All users of dropwatch are advised to upgrade to these updated packages, which fix this bug.

5.55.2. RHBA-2012:0383 – dropwatch bug fix update

An updated dropwatch package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The dropwatch package contains a utility that provides packet monitoring services.

Bug Fix

BZ#684713

Previously, the dropwatch utility could terminate unexpectedly with a segmentation fault. The failure was caused by a double-free error which occurred while issuing the start and stop messages. This update removes the freeing function calls from the underlying code, which prevents the dropwatch utility from crashing.

All users of dropwatch are advised to upgrade to this updated package, which fixes this bug.

5.56. DVD+RW-TOOLS

5.56.1. RHBA-2012:1320 – dvd+rw-tools bug fix update

Updated dvd+rw-tools packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The dvd+rw-tools packages contain a collection of tools to master DVD+RW/+R media.

Bug Fix

BZ#807474

Prior to this update, the growisofs utility wrote chunks of 32KB and reported an error during the last chunk when burning ISO image files that were not aligned to 32KB. This update allows the written chunk to be smaller than a multiple of 16 blocks.

All users of dvd+rw-tools are advised to upgrade to these updated packages, which fix this bug.

5.57. E2FSPROGS

5.57.1. RHBA-2012:0944 – e2fsprogs bug fix update

Updated e2fsprogs packages that fix two 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 second (ext2), third (ext3), and fourth (ext4) extended file systems.

Bug Fixes

BZ#786021

Prior to this update, checksums for backup group descriptors appeared to be wrong when the "e2fsck -b" option read these group descriptors and cleared UNINIT flags to ensure that all inodes were scanned. As a consequence, warning messages were sent during the process. This update recomputes checksums after the flags are changed. Now, "e2fsck -b" completes without these checksum warnings.

BZ#795846

Prior to this update, e2fsck could discard valid inodes when using the "-E discard" option. As a consequence, the file system could become corrupted. This update modifies the underlying code so that disk regions containing valid inodes are no longer discarded.

All users of e2fsprogs are advised to upgrade to these updated packages, which fix these bugs.

5.58. EFIBOOTMGR

5.58.1. RHBA-2012:0893 – efibootmgr bug fix update

An updated efibootmgr package that fixes one bug is 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#715216

In a Coverity Scan analysis, an allocation, which was not checked for errors, was discovered. With this update, the allocation is now checked for errors, thus the bug is fixed.

All users of efibootmgr are advised to upgrade to this updated package, which fixes this bug.

5.59. ELINKS

5.59.1. RHSA-2013:0250 – Moderate: elinks security update

An updated elinks package that fixes one security issue is now available for Red Hat Enterprise Linux 5 and 6.

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

ELinks is a text-based web browser. ELinks does not display any images, but it does support frames, tables, and most other HTML tags.

Security Fix

CVE-2012-4545

It was found that ELinks performed client credentials delegation during the client-to-server GSS security mechanisms negotiation. A rogue server could use this flaw to obtain the client's credentials and impersonate that client to other servers that are using GSSAPI.

This issue was discovered by Marko Myllynen of Red Hat.

All ELinks users are advised to upgrade to this updated package, which contains a backported patch to resolve the issue.

5.60. ESPEAK

5.60.1. RHBA-2012:1118 – espeak bug fix update

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

The espeak packages contain a software speech synthesizer for English and other languages. eSpeak uses a "formant synthesis" method, which allows many languages to be provided in a small size.

Bug Fix

BZ#789997

Previously, eSpeak manipulated the system sound volume. As a consequence, eSpeak could set the sound volume to maximum regardless of the amplitude specified. The sound volume management code has been removed from eSpeak, and now only PulseAudio manages the sound volume.

All users of espeak are advised to upgrade to these updated packages, which fix this bug.

5.61. EXPECT

5.61.1. RHBA-2012:0456 – expect bug fix update

An updated expect package that fixes multiple bugs is now available for Red Hat Enterprise Linux 6.

The "expect" package contains a tool for automating and testing interactive command line programs and Tk applications. Tcl is a portable and widely used scripting language, while Tk is a graphical toolkit that eases development of text-based and GUI applications.

Bug Fixes

BZ#674866

Prior to this update, the expect(1) manual page was not formatted properly. As a result, the content of the manual page was not readable. The formatting has been corrected to ensure easy readability.

BZ#735962

Prior to this update, the passmass script did not call the "su" binary with the full path (/bin/su). The passmass script has been modified to call "/bin/su" rather than "su", which is more secure.

BZ#742911

Due to incorrect characters matching, applications created by the autoexpect utility could terminate unexpectedly with a segmentation fault. With this update, the number of characters is matched correctly and applications created by autoexpect run successfully.

BZ#782859

Previously, the expect-devel subpackage contained a symbolic link to the expect library, which led to an unnecessary dependency. With this update, the link is located in the expect package.

All users of expect are advised to upgrade to this updated package, which fixes these bugs.

5.62. FCOE-TARGET-UTILS

5.62.1. RHBA-2012:0854 – fcoe-target-utils bug fix and enhancement update

Updated fcoe-target-utils packages that fix three bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The fcoe-target-utils packages provide a command line interface for configuring FCoE LUNs (Fibre Channel over Ethernet Logical Unit Numbers) and backstores.

Bug Fixes

BZ#752699

Prior to this update, starting targetadmin without the fcoe-target utility could cause the following output:

```
OSError: [Errno 2] No such file or directory: '/sys/kernel/config/target
```

This update modifies the underlying code so that now a warning message is displayed if targetcli is invoked without running the fcoe-target service.

BZ#813664

Prior to this update, fcoe-target-utils used the executable name "targetadmin" which did not reflect the current name in the upstream version. This update changes the name to "targetcli", to match the upstream version.

BZ#815981

Prior to this update, the configuration state was saved to "tcm_start.sh", and the fcoe-target init script restored the state from this file when the fcoe-target service was started. For increased reliability, this update uses a new method to save and restore fcoe-target configuration; it is now saved to "/etc/target/saveconfig.json".

Enhancement

BZ#750277

Prior to this update, the fcoe-target-utils packages for the Fibre Channel over Ethernet (FCoE) target mode were available only as technical preview. With this update, the fcoe-target-utils packages are fully supported in Red Hat Enterprise Linux 6.

All users of fcoe-target-utils are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.63. FCOE-UTILS

5.63.1. RHBA-2012:0851 – fcoe-utils bug fix and enhancement update

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

The fcoe-utils package allows users to use Fibre Channel over Ethernet (FCoE). The package contains the fcoeadm command-line tool for configuring FCoE interfaces, and the fcoemon service to configure DCB (Data Center Bridging) Ethernet QoS filters.

The fcoe-utils package has been upgraded to upstream version 1.0.22, which provides a number of bug fixes and enhancements over the previous version, including bash-completion enhancements and changing of the default FCoE interface names from 'device.vlan-fcoe' to 'device.vlan'. The -f (--fipvlan) option can be used to apply the previous behavior. (BZ#788511)

Bug Fix

BZ#804936

The "service fcoe status" command returned an incorrect return value when the fcoe service was running. With this update, the underlying code has been modified and fcoe now returns the correct code under these circumstances.

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

5.64. FEBOOTSTRAP

5.64.1. RHEA-2012:0775 – febootstrap bug fix and enhancement update

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

The febootstrap packages provide a tool to create a basic Red Hat Enterprise Linux or Fedora filesystem, and builds initramfs (initrd.img) or filesystem images.

The febootstrap packages have been upgraded to upstream version 3.12, which provides a number of bug fixes and enhancements over the previous version. (BZ#719877)

All febootstrap users are advised upgrade to these updated packages, which fix these bugs and add these enhancements.

5.65. FENCE-AGENTS

5.65.1. RHBA-2012:1439 – fence-agents bug fix update

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

The fence-agents packages provide the Red Hat fence agents to handle remote power management for cluster devices. The fence-agents allow failed or unreachable nodes to be forcibly restarted and removed from the cluster.

Bug Fix

BZ#872620

The speed of fencing is critical because otherwise, broken nodes have more time to corrupt data. Prior to this update, the operation of the fence_vmware_soap fencing agent was slow and could corrupt data when used on the VMWare vSphere platform with hundreds of virtual machines. This

update fixes a problem with virtual machines that do not have a valid UUID, which can be created during failed P2V (Physical-to-Virtual) processes. Now, the fencing process is also much faster and it does not terminate if a virtual machines without an UUID is encountered.

All users of fence-agents are advised to upgrade to these updated packages, which fix this bug.

5.65.2. RHBA-2012:0943 – fence-agents bug fix and enhancement update

Updated fence-agents packages that fix various bugs and add an enhancement are now available for Red Hat Enterprise Linux 6.

The fence-agents package contains a collection of scripts to handle remote power management for cluster devices. They allow failed or unreachable cluster nodes to be forcibly restarted and removed from the cluster.

Bug Fixes

BZ#769681

The fence_rhevm fencing agent uses the Red Hat Enterprise Virtualization API to check the power status ("on" or "off") of a virtual machine. In addition to the "up" and "down" states, the API includes number of other states. Previously, only if the machine was in the "up" state, the "on" power status was returned. The "off" status was returned for all other states even if the machine was running. This allowed for successful fencing before the machine was really powered off. With this update, the fence_rhevm agent detects the power status of a cluster node more conservatively, and the "off" status is returned only if the machine is actually powered off, that is in the "down" state.

BZ#772597

Previously, the fence_soap_vmware fence agent was not able to work with more than one hundred machines in a cluster. Consequently, fencing a cluster node running in a virtual machine on VMWare with the fence_soap_vmware fence agent failed with the "KeyError: 'config.uuid'" error message. With this update, the underlying code has been fixed to support fencing on such clusters.

BZ#740484

Previously, the fence_ipmilan agent failed to handle passwd_script argument values that contained space characters. Consequently, it was impossible to use a password script that required additional parameters. This update ensures that fence_ipmilan accepts and properly parses values for the passwd_script argument with spaces.

BZ#771211

Previously, the fence_vmware_soap fence agent did not expose the proper virtual machine path for fencing. With this update, fence_vmware_soap has been fixed to support this virtual machine identification.

BZ#714841

Previously, certain fence agents did not generate correct metadata output. As a result, it was not possible to use the metadata for automatic generation of manual pages and user interfaces. With this update, all fence agents generate their metadata as expected.

BZ#771936

Possible buffer overflow and null dereference defects were found by automatic tools. With this update, these problems have been fixed.

BZ#785091

Fence agents that use an identity file for SSH terminated unexpectedly when a password was expected but was not provided. This bug has been fixed and proper error messages are returned in the described scenario.

BZ#787706

The fence_ipmilan fence agent did not respect the power_wait option and did not wait after sending the power-off signal to a device. Consequently, the device could terminate its shutdown sequence. This bug has been fixed and fence_ipmilan now waits before shutting down a machine as expected.

BZ#741339

The fence_scsi agent creates the fence_scsi.dev file that contains a list of devices that the node registered with during an unfence operation. This file was unlinked for every unfence action. Consequently, if multiple fence device entries were used in the cluster.conf file, fence_scsi.dev only contained the devices that the node registered with during the most recent unfence action. Now, instead of the unlink call, if the device currently being registered does not exists in fence_scsi.dev, it is added to the file.

BZ#804169

If the "delay" option was set to more than 5 seconds while a fence device was connected via the telnet_ssl utility, the connection timed out and the fence device failed. Now, the "delay" option is applied before the connection is opened, thus fixing this bug.

BZ#806883

Previously, XML metadata returned by a fence agent incorrectly listed all attributes as "unique". This update fixes this problem and the attributes are now marked as unique only when this information is valid.

BZ#806912

This update fixes a typographical error in an error message in the fence_ipmilan agent.

BZ#806897

Prior to this update, the fence agent for IPMI (Intelligent Platform Management Interface) could return an invalid return code when the "-M cycle" option was used. This invalid return code could cause invalid interpretation of a fence action, eventually causing the cluster to become unresponsive. This bug has been fixed and only predefined return codes are now returned in the described scenario.

BZ#804805

Previously, the fence_brocade fence agent did not distinguish the "action" option from the standard "option" option. Consequently, the "action" option was ignored and the node was always fenced. This bug has been fixed and both options are now properly recognized and acted upon.

Enhancement

BZ#742003

This updates adds the feature to access Fujitsu RSB fencing device using secure shell.

Users of fence-agents are advised to upgrade to these updated packages, which fix these bugs and add this enhancements.

5.66. FENCE-VIRT

5.66.1. RHBA-2012:0800 – fence-virt bug fix update

Updated fence-virt packages that fix four 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#753974

Prior to this update, the libvirt-qpid plug-in did not handle exceptions correctly. As a consequence, the fence_virtd daemon could unexpectedly terminate with a segmentation fault if the connection to the specified qpid daemon failed. This update modifies the exception handling. Now, the fencing operation works as expected.

BZ#758392

Prior to this update, the hashing utility sha_verify did not handle errors correctly when a key file could not be read. As a consequence, the fence_virtd daemon could unexpectedly terminate with a segmentation fault when receiving a fencing request if fence_virtd failed to read the specified key file during startup. This update modifies the error handling if a key file cannot be read. Now, fence_virtd no longer terminates under these conditions.

BZ#761215

Prior to this update, the XML example for serial mode in the fence_virt.conf(5) man page contained an incorrect closing tag. This update corrects this tag.

BZ#806949

Prior to this update, the libvirt-qpid plug-in was linked directly against the qpid libraries instead of only the qmfv2 library. As a consequence, newer versions of the qpid libraries could not be used with the libvirt-qpid plug-in. This update no longer links against the qpid libraries directly. Now, also newer qpid libraries can be used with libvirt-qpid.

BZ#809101

Prior to this update, the fence_virtd.conf manpage and the fence_virtd.conf generator incorrectly stated that by default, fence_virtd listened on all network interfaces. Both have been amended to state that by default, fence_virtd listens on the default network interface.

All users of fence-virt are advised to upgrade to these updated packages, which fix these bugs.

5.67. FILE

5.67.1. RHBA-2012:1339 – file bug fix update

Updated file packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The "file" command is used to identify a particular file according to the type of data contained in the file. The command can identify various file types, including ELF binaries, system libraries, RPM packages, and different graphics formats.

Bug Fixes

BZ#795425

The file utility did not contain a "magic" pattern for detecting QED images and was therefore not able to detect such images. A new "magic" pattern for detecting QED images has been added, and the file utility now detects these images as expected.

BZ#795761

The file utility did not contain a "magic" pattern for detecting VDI images and was therefore not able to detect such images. A new "magic" pattern for detecting VDI images has been added, and the file utility now detects these images as expected.

BZ#797784

Previously, the file utility did not attempt to load "magic" patterns from the ~/.magic.mgc file, which caused "magic" patterns stored in this file to be unusable. This update modifies the file utility so it now attempts to load the ~/.magic.mgc file. The file is loaded if it exists and "magic" patterns defined in this file work as expected.

BZ#801711

Previously, the file utility used read timeout when decompressing files using the "-z" option. As a consequence, the utility was not able to detect files compressed by the bzip2 tool. The underlying source code has been modified so that file no longer uses read timeout when decompressing compressed files. Compressed files are now detected as expected when using the "-z" option.

BZ#859834

Previously, the file utility contained multiple "magic" patterns to detect output of the "dump" backup tool. On big-endian architectures, the less detailed "magic" pattern was used and output of the file utility was inconsistent. The less detailed "magic" pattern has been removed, and only one, more detailed, "magic" pattern to detect "dump" output is used now.

All users of file are advised to upgrade to these updated packages, which fix these bugs.

5.67.2. RHBA-2012:0391 – file bug fix update

Updated file packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The file command is used to identify a particular file according to the type of data contained in the file. The command can identify various file types, including ELF binaries, system libraries, RPM packages, and different graphics formats.

Bug Fixes

BZ#688136

Previously, the file utility contained "magic" patterns that incorrectly detected files according to one byte only. Unicode text files starting with that particular byte could be therefore incorrectly recognized as DOS executable files. This update removes the problematic patterns. Patterns that match less than 16 bits are no longer accepted, and the utility no longer detects Unicode files as DOS executables.

BZ#709846

Previously, the "magic" pattern for detection of Dell BIOS headers was outdated. As a consequence, the file utility did not detect newer BIOS formats. The "magic" pattern has been updated, and the file utility now detects new formats of Dell BIOS properly.

BZ#719583

Previously, users were allowed to add new "magic" files only into the home directory. As a consequence, users were not able to configure "magic" patterns for certain special file formats system-wide. With this update, a backported patch provides a way to read "magic" patterns from the /etc/magic file.

BZ#733229

Previously, "magic" patterns for Python were insufficient. The file utility was therefore unable to detect a Python script according to the Python function definition. With this update, detection of Python is improved, and Python scripts are properly recognized.

BZ#747999

Previously, the file utility did not contain a "magic" pattern for detection of files compressed using the LZMA algorithm. As a consequence, the file utility was unable to detect these files. This update adds the missing "magic" pattern, and LZMA compressed files are now detected as expected.

BZ#758109

Previously, the file utility did not contain a "magic" pattern to detect the swap signature on Itanium microprocessors. As a consequence, the file utility was unable to detect the signature. This update adds the missing "magic" pattern, and the swap signature on Itanium microprocessors is detected as expected.

BZ#760083

Previously, the file utility did not parse the name of an RPM package from the RPM file. As a consequence, the utility did not print the name of the RPM package. This update adds a "magic" pattern for RPM package name parsing, and the name is now printed as expected.

All users of file are advised to upgrade to these updated packages, which fix these bugs.

5.68. FIREFOX

5.68.1. RHSA-2013:0271 – Critical: firefox security update

Updated firefox packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fixes

CVE-2013-0775, CVE-2013-0780, CVE-2013-0782, CVE-2013-0783

Several flaws were found in the processing of malformed web content. A web page containing malicious content could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2013-0776

It was found that, after canceling a proxy server's authentication prompt, the address bar continued to show the requested site's address. An attacker could use this flaw to conduct phishing attacks by tricking a user into believing they are viewing a trusted site.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Nils, Abhishek Arya, Olli Pettay, Christoph Diehl, Gary Kwong, Jesse Ruderman, Andrew McCreight, Joe Drew, Wayne Mery, and Michal Zalewski as the original reporters of these issues.

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 17.0.3 ESR:

http://www.mozilla.org/security/known-vulnerabilities/firefoxESR.html

Note that due to a Kerberos credentials change, the following configuration steps may be required when using Firefox 17.0.3 ESR with the Enterprise Identity Management (IPA) web interface:

https://access.redhat.com/site/solutions/294303



IMPORTANT

Firefox 17 is not completely backwards-compatible with all Mozilla add-ons and Firefox plug-ins that worked with Firefox 10.0. Firefox 17 checks compatibility on first-launch, and, depending on the individual configuration and the installed add-ons and plug-ins, may disable said Add-ons and plug-ins, or attempt to check for updates and upgrade them. Add-ons and plug-ins may have to be manually updated.

All Firefox users should upgrade to these updated packages, which contain Firefox version 17.0.3 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.68.2. RHSA-2012:1088 – Critical: firefox security update

Updated firefox packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fixes

CVE-2012-1948, CVE-2012-1951, CVE-2012-1952, CVE-2012-1953, CVE-2012-1954, CVE-2012-1958, CVE-2012-1962, CVE-2012-1967

A web page containing malicious content could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-1959

A malicious web page could bypass same-compartment security wrappers (SCSW) and execute arbitrary code with chrome privileges.

CVE-2012-1966

A flaw in the context menu functionality in Firefox could allow a malicious website to bypass intended restrictions and allow a cross-site scripting attack.

CVE-2012-1950

A page different to that in the address bar could be displayed when dragging and dropping to the address bar, possibly making it easier for a malicious site or user to perform a phishing attack.

CVE-2012-1955

A flaw in the way Firefox called history.forward and history.back could allow an attacker to conceal a malicious URL, possibly tricking a user into believing they are viewing a trusted site.

CVE-2012-1957

A flaw in a parser utility class used by Firefox to parse feeds (such as RSS) could allow an attacker to execute arbitrary JavaScript with the privileges of the user running Firefox. This issue could have affected other browser components or add-ons that assume the class returns sanitized input.

CVE-2012-1961

A flaw in the way Firefox handled X-Frame-Options headers could allow a malicious website to perform a clickjacking attack.

CVE-2012-1963

A flaw in the way Content Security Policy (CSP) reports were generated by Firefox could allow a malicious web page to steal a victim's OAuth 2.0 access tokens and OpenID credentials.

CVE-2012-1964

A flaw in the way Firefox handled certificate warnings could allow a man-in-the-middle attacker to create a crafted warning, possibly tricking a user into accepting an arbitrary certificate as trusted.

CVE-2012-1965

A flaw in the way Firefox handled feed:javascript URLs could allow output filtering to be bypassed, possibly leading to a cross-site scripting attack.

The nss update RHBA-2012:0337 for Red Hat Enterprise Linux 5 and 6 introduced a mitigation for the CVE-2011-3389 flaw. For compatibility reasons, it remains disabled by default in the nss packages. This update makes Firefox enable the mitigation by default. It can be disabled by setting the NSS_SSL_CBC_RANDOM_IV environment variable to 0 before launching Firefox. (BZ#838879)

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 10.0.6 ESR.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Benoit Jacob, Jesse Ruderman, Christian Holler, Bill McCloskey, Abhishek Arya, Arthur Gerkis, Bill Keese, moz_bug_r_a4, Bobby Holley, Code Audit Labs, Mariusz Mlynski, Mario Heiderich, Frédéric Buclin, Karthikeyan Bhargavan, Matt McCutchen, Mario Gomes, and Soroush Dalili as the original reporters of these issues.

All Firefox users should upgrade to these updated packages, which contain Firefox version 10.0.6 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.68.3. RHSA-2012:1210 – Critical: firefox security update

Updated firefox packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fixes

CVE-2012-1970, CVE-2012-1972, CVE-2012-1973, CVE-2012-1974, CVE-2012-1975, CVE-2012-1976, CVE-2012-3956, CVE-2012-3957, CVE-2012-3958, CVE-2012-3959, CVE-2012-3960, CVE-2012-3961, CVE-2012-3962, CVE-2012-3963, CVE-2012-3964

A web page containing malicious content could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-3969, CVE-2012-3970

A web page containing a malicious Scalable Vector Graphics (SVG) image file could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-3967, CVE-2012-3968

Two flaws were found in the way Firefox rendered certain images using WebGL. A web page containing malicious content could cause Firefox to crash or, under certain conditions, possibly execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-3966

A flaw was found in the way Firefox decoded embedded bitmap images in Icon Format (ICO) files. A web page containing a malicious ICO file could cause Firefox to crash or, under certain conditions, possibly execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-3980

A flaw was found in the way the "eval" command was handled by the Firefox Web Console. Running "eval" in the Web Console while viewing a web page containing malicious content could possibly cause Firefox to execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-3972

An out-of-bounds memory read flaw was found in the way Firefox used the format-number feature of XSLT (Extensible Stylesheet Language Transformations). A web page containing malicious content could possibly cause an information leak, or cause Firefox to crash.

CVE-2012-3976

It was found that the SSL certificate information for a previously visited site could be displayed in the address bar while the main window displayed a new page. This could lead to phishing attacks as attackers could use this flaw to trick users into believing they are viewing a trusted site.

CVE-2012-3978

A flaw was found in the location object implementation in Firefox. Malicious content could use this flaw to possibly allow restricted content to be loaded.

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 10.0.7 ESR.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Gary Kwong, Christian Holler, Jesse Ruderman, John Schoenick, Vladimir Vukicevic, Daniel Holbert, Abhishek Arya, Frédéric Hoguin, miaubiz, Arthur Gerkis, Nicolas Grégoire, Mark Poticha, moz_bug_r_a4, and Colby Russell as the original reporters of these issues.

All Firefox users should upgrade to these updated packages, which contain Firefox version 10.0.7 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.68.4. RHSA-2012:1407 – Critical: firefox security update

Updated firefox packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fix

CVE-2012-4194, CVE-2012-4195, CVE-2012-4196

Multiple flaws were found in the location object implementation in Firefox. Malicious content could be used to perform cross-site scripting attacks, bypass the same-origin policy, or cause Firefox to execute arbitrary code.

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 10.0.10 ESR:

http://www.mozilla.org/security/known-vulnerabilities/firefoxESR.html

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Mariusz Mlynski, moz_bug_r_a4, and Antoine Delignat-Lavaud as the original reporters of these issues.

All Firefox users should upgrade to these updated packages, which contain Firefox version 10.0.10 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.68.5. RHSA-2013:0144 – Critical: firefox security update

Updated firefox packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fixes

CVE-2013-0744, CVE-2013-0746, CVE-2013-0750, CVE-2013-0753, CVE-2013-0754, CVE-2013-0762, CVE-2013-0766, CVE-2013-0767, CVE-2013-0769

Several flaws were found in the processing of malformed web content. A web page containing malicious content could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2013-0758

A flaw was found in the way Chrome Object Wrappers were implemented. Malicious content could be used to cause Firefox to execute arbitrary code via plug-ins installed in Firefox.

CVE-2013-0759

A flaw in the way Firefox displayed URL values in the address bar could allow a malicious site or user to perform a phishing attack.

CVE-2013-0748

An information disclosure flaw was found in the way certain JavaScript functions were implemented in Firefox. An attacker could use this flaw to bypass Address Space Layout Randomization (ASLR) and other security restrictions.

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 10.0.12 ESR:

http://www.mozilla.org/security/known-vulnerabilities/firefoxESR.html

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Atte Kettunen, Boris Zbarsky, pa_kt, regenrecht, Abhishek Arya, Christoph Diehl, Christian Holler, Mats Palmgren, Chiaki Ishikawa, Mariusz Mlynski, Masato Kinugawa, and Jesse Ruderman as the original reporters of these issues.

All Firefox users should upgrade to these updated packages, which contain Firefox version 10.0.12 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.68.6. RHSA-2012:1350 – Critical: firefox security and bug fix update

Updated firefox packages that fix several security issues and one bug are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fixes

CVE-2012-3982, **CVE-2012-3988**, **CVE-2012-3990**, **CVE-2012-3995**, **CVE-2012-4179**, **CVE-2012-4180**, **CVE-2012-4181**, **CVE-2012-4182**, **CVE-2012-4183**, **CVE-2012-4185**, **CVE-2012-4186**, **CVE-2012-4187**, **CVE-2012-4188**

Several flaws were found in the processing of malformed web content. A web page containing malicious content could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-3986, **CVE-2012-3991**

Two flaws in Firefox could allow a malicious website to bypass intended restrictions, possibly leading to information disclosure, or Firefox executing arbitrary code. Note that the information disclosure issue could possibly be combined with other flaws to achieve arbitrary code execution.

CVE-2012-1956, **CVE-2012-3992**, **CVE-2012-3994**

Multiple flaws were found in the location object implementation in Firefox. Malicious content could be used to perform cross-site scripting attacks, script injection, or spoofing attacks.

CVE-2012-3993, **CVE-2012-4184**

Two flaws were found in the way Chrome Object Wrappers were implemented. Malicious content could be used to perform cross-site scripting attacks or cause Firefox to execute arbitrary code.

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 10.0.8 ESR.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Christian Holler, Jesse Ruderman, Soroush Dalili, miaubiz, Abhishek Arya, Atte Kettunen, Johnny Stenback, Alice White, moz_bug_r_a4, and Mariusz Mlynski as the original reporters of these issues.

Bug Fix

BZ#809571, **BZ#816234**

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" (without quotes) 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.

All Firefox users should upgrade to these updated packages, which contain Firefox version 10.0.8 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.68.7. RHSA-2012:1482 – Critical: firefox security update

Updated firefox packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Firefox is an open source web browser. XULRunner provides the XUL Runtime environment for Mozilla Firefox.

Security Fixes

CVE-2012-4214, CVE-2012-4215, CVE-2012-4216, CVE-2012-5829, CVE-2012-5830, CVE-2012-5833, CVE-2012-5835, CVE-2012-5839, CVE-2012-5840, CVE-2012-5842

Several flaws were found in the processing of malformed web content. A web page containing malicious content could cause Firefox to crash or, potentially, execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-4202

A buffer overflow flaw was found in the way Firefox handled GIF (Graphics Interchange Format) images. A web page containing a malicious GIF image could cause Firefox to crash or, possibly, execute arbitrary code with the privileges of the user running Firefox.

CVE-2012-4210

A flaw was found in the way the Style Inspector tool in Firefox handled certain Cascading Style Sheets (CSS). Running the tool (Tools -> Web Developer -> Inspect) on malicious CSS could result in the execution of HTML and CSS content with chrome privileges.

CVE-2012-4207

A flaw was found in the way Firefox decoded the HZ-GB-2312 character encoding. A web page containing malicious content could cause Firefox to run JavaScript code with the permissions of a different website.

CVE-2012-4209

A flaw was found in the location object implementation in Firefox. Malicious content could possibly use this flaw to allow restricted content to be loaded by plug-ins.

CVE-2012-5841

A flaw was found in the way cross-origin wrappers were implemented. Malicious content could use this flaw to perform cross-site scripting attacks.

CVE-2012-4201

A flaw was found in the evalInSandbox implementation in Firefox. Malicious content could use this flaw to perform cross-site scripting attacks.

For technical details regarding these flaws, refer to the Mozilla security advisories for Firefox 10.0.11 ESR:

http://www.mozilla.org/security/known-vulnerabilities/firefoxESR.html

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Abhishek Arya, miaubiz, Jesse Ruderman, Andrew McCreight, Bob Clary, Kyle Huey, Atte Kettunen, Mariusz Mlynski, Masato Kinugawa, Bobby Holley, and moz_bug_r_a4 as the original reporters of these issues.

All Firefox users should upgrade to these updated packages, which contain Firefox version 10.0.11 ESR, which corrects these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.69. FIRSTBOOT

5.69.1. RHEA-2012:0928 — firstboot enhancement update

Updated firstboot packages that add two enhancements are now available for Red Hat Enterprise Linux 6.

The firstboot utility runs after installation and guides the user through a series of steps that allows for easier configuration of the machine.

Enhancements

BZ#704187

Prior to this update, the firstboot utility did not allow users to change the timezone. This update adds the timezone module to firstboot so that users can now change the timezone in the reconfiguration mode.

BZ#753658

Prior to this update, the firstboot service did not provide a status option. This update adds the "firstboot service status" option to show if firstboot is scheduled to run on the next boot or not.

All users of firstboot are advised to upgrade to these updated packages, which add these enhancements.

5.70. FLASH-PLUGIN

5.70.1. RHSA-2012:1173 — Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes one security issue is now available for Red Hat Enterprise Linux 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fix

CVE-2012-1535

This update fixes one vulnerability in Adobe Flash Player. This vulnerability is detailed on the Adobe security page APSB12-18. Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.238.

5.70.2. RHSA-2012:1431 – Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fix

CVE-2012-5274, CVE-2012-5275, CVE-2012-5276, CVE-2012-5277, CVE-2012-5278, CVE-2012-5279, CVE-2012-5280

This update fixes several vulnerabilities in Adobe Flash Player. These vulnerabilities are detailed in the Adobe Security bulletin APSB12-24 . Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.251.

5.70.3. RHSA-2013:0243 – Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes two security issues is now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fix

CVE-2013-0633, CVE-2013-0634

This update fixes two vulnerabilities in Adobe Flash Player. These vulnerabilities are detailed in the Adobe Security bulletin APSB13-04 . Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.262.

5.70.4. RHSA-2012:1346 – Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fix

CVE-2012-5248, CVE-2012-5249, CVE-2012-5250, CVE-2012-5251, CVE-2012-5252, CVE-2012-5253, CVE-2012-5254, CVE-2012-5255, CVE-2012-5256, CVE-2012-5257, CVE-2012-5258, CVE-2012-5259, CVE-2012-5260, CVE-2012-5261, CVE-2012-5262, CVE-2012-5263, CVE-2012-5264, CVE-2012-5265, CVE-2012-5266, CVE-2012-5267, CVE-2012-5268, CVE-2012-5269, CVE-2012-5270, CVE-2012-5271, CVE-2012-5272

This update fixes several vulnerabilities in Adobe Flash Player. These vulnerabilities are detailed on the Adobe security page APSB12-22 . Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.243.

5.70.5. RHSA-2012:1569 – Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes three security issues is now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fix

CVE-2012-5676, CVE-2012-5677, CVE-2012-5678

This update fixes three vulnerabilities in Adobe Flash Player. These vulnerabilities are detailed in the Adobe Security bulletin APSB12-27 . Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.258.

5.70.6. RHSA-2013:0149 – Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes one security issue is now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fix

CVE-2013-0630

This update fixes one vulnerability in Adobe Flash Player. This vulnerability is detailed in the Adobe Security bulletin APSB13-01. Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.261.

5.70.7. RHSA-2013:0254 – Critical: flash-plugin security update

An updated Adobe Flash Player package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

Security Fixes

CVE-2013-0638, CVE-2013-0639, CVE-2013-0642, CVE-2013-0644, CVE-2013-0645, CVE-2013-0647, CVE-2013-0649, CVE-2013-1365, CVE-2013-1366, CVE-2013-1367, CVE-2013-1368, CVE-2013-1369, CVE-2013-1370, CVE-2013-1372, CVE-2013-1373, CVE-2013-1374

This update fixes several vulnerabilities in Adobe Flash Player. These vulnerabilities are detailed in the Adobe Security bulletin APSB13-05 . Specially-crafted SWF content could cause flash-plugin to crash or, potentially, execute arbitrary code when a victim loads a page containing the malicious SWF content.

CVE-2013-0637

A flaw in flash-plugin could allow an attacker to obtain sensitive information if a victim were tricked into visiting a specially-crafted web page.

All users of Adobe Flash Player should install this updated package, which upgrades Flash Player to version 11.2.202.270.

5.70.8. RHEA-2012:0980 – flash-plugin enhancement update

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

The flash-plugin package contains a Mozilla Firefox compatible Adobe Flash Player web browser plug-in.

The Adobe Flash Player web browser plug-in has been upgraded to upstream version 11.2.202.236, which provides a number of enhancements over the previous version. (BZ#800030)

All users of Adobe Flash Player are advised to upgrade these updated packages, which add these enhancements.

5.71. FONTFORGE

5.71.1. RHBA-2012:0351 – fontforge bug fix update

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

FontForge is a font editor for outline and bitmap fonts. FontForge supports a range of font formats, including PostScript, TrueType, OpenType and CID-keyed fonts.

Bug Fix

BZ#676607

Previously, the "configure.in" file did not include information on how to handle 64-bit PowerPC architectures. Attempting to install the fontforge-devel multilib PowerPC and 64-PowerPC RPM packages on the same 64-bit PowerPC machine led to conflicts between those packages. This update modifies the "configure.in" file, so that fontforge-devel multilib RPM packages are allowed to be installed on the same machine. The conflicts no longer occur in the described scenario.

All users of fontforge are advised to upgrade to these updated packages, which fix this bug.

5.72. FPRINTD

5.72.1. RHBA-2012:0912 – fprintd bug fix update

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

The fprintd packages contains a D-Bus service to access fingerprint readers.

Bug Fix

BZ#665837

Previously, if no USB support was available on a machine (for example, virtual machines on a hypervisor that disabled USB support for guests), the fprintd daemon received the SIGABRT signal, and therefore terminated abnormally. Such crashes did not cause any system failure; however, the Automatic Bug Reporting Tool (ABRT) was alerted every time. With this update, the underlying code has been modified so that the fprintd daemon now exits gracefully on machines with no USB support.

All users of fprintd are advised to upgrade to these updated packages, which fix this bug.

5.73. FREERADIUS

5.73.1. RHSA-2012:1326 – Moderate: freeradius security update

Updated freeradius packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

FreeRADIUS is a high-performance and highly configurable free Remote Authentication Dial In User Service (RADIUS) server, designed to allow centralized authentication and authorization for a network.

Security Fix

CVE-2012-3547

A buffer overflow flaw was discovered in the way radiusd handled the expiration date field in X.509 client certificates. A remote attacker could possibly use this flaw to crash radiusd if it were configured to use the certificate or TLS tunnelled authentication methods (such as EAP-TLS, EAP-TTLS, and PEAP).

Red Hat would like to thank Timo Warns of PRESENSE Technologies GmbH for reporting this issue.

Users of FreeRADIUS are advised to upgrade to these updated packages, which contain a backported patch to correct this issue. After installing the update, radiusd will be restarted automatically.

5.73.2. RHBA-2012:0881 – freeradius bug fix and enhancement update

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

FreeRADIUS is an open-source Remote Authentication Dial In User Service (RADIUS) server which allows RADIUS clients to perform authentication against the RADIUS server. The RADIUS server may optionally perform accounting of its operations using the RADIUS protocol.

The freeradius packages have been upgraded to upstream version 2.1.12, which provides a number of bug fixes and enhancements over the previous version. (BZ#736878)

Bug Fixes

BZ#787116

The radtest command-line argument to request the PPP hint option was not parsed correctly. Consequently, radclient did not add the PPP hint to the request packet and the test failed. This update corrects the problem and radtest now functions as expected.

BZ#705723

After log rotation, the freeradius logrotate script failed to reload the radiusd daemon after a log rotation and log messages were lost. This update has added a command to the freeradius logrotate script to reload the radiusd daemon and the radiusd daemon reinitializes and reopens its log files after log rotation as expected.

BZ#712803

The radtest argument with the eap-md5 option failed because it passed the IP family argument when invoking the radeapclient utility and the radeapclient utility did not recognize the IP family. The radeapclient now recognizes the IP family argument and radtest now works with eap-md5 as expected.

BZ#700870

Previously, freeradius was compiled without the "--with-udpfromto" option. Consequently, with a multihomed server and explicitly specifying the IP address, freeradius sent the reply from the wrong IP address. With this update, freeradius has been built with the --with-udpfromto configuration option and the RADIUS reply is always sourced from the IP the request was sent to.

BZ#753764

The password expiration field for local passwords was not checked by the unix module and the debug information was erroneous. Consequently, a user with an expired password in the local password file was authenticated despite having an expired password. With this update, check of the password expiration has been modified. A user with an expired local password is denied access and correct debugging information is written to the log file.

BZ#690756

Due to invalid syntax in the PostgreSQL admin schema file, the FreeRADIUS PostgreSQL tables failed to be created. With this update, the syntax has been adjusted and the tables are created as expected.

BZ#782905

When FreeRADIUS received a request, it sometimes failed with the following message:

```
WARNING: Internal sanity check failed in event handler for request 6
```

This bug was fixed by upgrading to upstream version 2.1.12.

BZ#810605

FreeRADIUS has a thread pool that will dynamically grow based on load. If multiple threads using the rlm_perl() function are spawned in quick succession, freeradius sometimes terminated unexpectedly with a segmentation fault due to parallel calls to the rlm_perl_clone() function. With this update, mutex for the threads has been added and the problem no longer occurs.

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

5.74. FREETYPE

5.74.1. RHSA-2013:0216 – Important: freetype security update

Updated freetype packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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(s) associated with each description below.

FreeType is a free, high-quality, portable font engine that can open and manage font files. It also loads, hints, and renders individual glyphs efficiently.

Security Fix

CVE-2012-5669

A flaw was found in the way the FreeType font rendering engine processed certain Glyph Bitmap Distribution Format (BDF) fonts. If a user loaded a specially-crafted font file with an application linked against FreeType, it could cause the application to crash or, possibly, execute arbitrary code with the privileges of the user running the application.

Users are advised to upgrade to these updated packages, which contain a backported patch to correct this issue. The X server must be restarted (log out, then log back in) for this update to take effect.

5.75. FTP

5.75.1. RHBA-2012:1192 – ftp bug fix update

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

The ftp package provides 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#783868

Prior to this update, using the ftp command "put" when the stack size was set to unlimited caused the sysconf(_SC_ARG_MAX) function to return -1, which in turn resulted in the malloc() function being called with an argument of 0 and causing an "Out of memory" message to be displayed. With this update, the underlying source code has been improved to allocate a reasonable minimum of memory. As a result, the "Out of memory" message no longer appears if the stack size was previously set to unlimited.

All users of ftp are advised to upgrade to these updated packages, which fix this bug.

5.75.2. RHBA-2012:1452 – ftp bug fix and enhancement update

Updated ftp packages that two bugs and add one enhancement 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 Fixes

BZ#871072

Previous implementation of FTP did not free the memory allocated for its commands correctly. Consequently, memory leaks occurred whenever the "append", "put" and "send" commands were run. With this update, the underlying source code has been corrected and allocated memory is now freed as expected.

BZ#871547

Previously, the size of the buffer used for an FTP macro definition was limited to 200 characters. Therefore, if the size of the macro was larger than 200 characters, the buffer overflowed and the FTP client terminated unexpectedly. This update extends the buffer of the FTP macro to match the size of the FTP command line limit, which is now 4296 characters. The FTP client no longer crashes in this scenario.

Enhancement

BZ#871060

Previously, the command line width in the FTP client was limited to 200 characters. With this update, the maximum possible length of the FTP command line has been extended to 4296 characters.

All users of ftp are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.75.3. RHBA-2012:1444 – 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#869858

Prior to this update, the ftp client could encounter a buffer overflow and aborted if a macro longer than 200 characters was defined and then used after a connection. This update modifies the underlying code and the buffer that holds memory for the macro name was extended. Now, ftp matches the length of the command line limit and the ftp client no longer aborts when a macro with a long name is executed.

All users of ftp are advised to upgrade to these updated packages, which fix this bug.

5.75.4. RHBA-2012:1354 – ftp bug fix update

Updated ftp packages that fix four bugs 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 Fixes

BZ#665337

Previously, the command line width in the ftp client was limited to 200 characters. With this update, the maximum possible length of the FTP command line is extended to 4296 characters.

BZ#786004

Prior to this update, "append", "put", and "send" commands were causing system memory to leak. The memory holding the ftp command was not freed appropriately. With this update, the underlying

source code has been improved to correctly free the system resources and the memory leaks are no longer present.

BZ#849940

Previously, the ftp client could not be invoked to run directly in the active mode. This functionality has been added to the source code and documented in the manual page. The client can now be executed with an additional "-A" command line parameter and will run in the active mode.

BZ#852636

Previously, the ftp client hung up when the ftp-data port (20) was not available (e.g. was blocked). The client then had to be terminated manually. Additional logic has been added to the source code. With this update, ftp has an internal timeout set to 30 seconds. If there is no answer from the server when this time has passed, ftp will now gracefully time out and not hang up.

All users of ftp are advised to upgrade to these updated packages, which fix these bugs.

5.76. GAWK

5.76.1. RHBA-2012:1146 – gawk bug fix update

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

The gawk packages provide the GNU version of the text processing utility awk. Awk interprets a special-purpose programming language to do quick and easy text pattern matching and reformatting jobs.

Bug Fix

BZ#829558

Prior to this update, the "re_string_skip_chars" function incorrectly used the character count instead of the raw length to estimate the string length. As a consequence, any text in multi-byte encoding that did not use the UTF-8 format failed to be processed correctly. This update modifies the underlying code so that the correct string length is used. multi-byte encoding is processed correctly.

All users of gawk requiring multi-byte encodings that do not use UTF-8 are advised to upgrade to these updated packages, which fix this bug.

5.76.2. RHBA-2012:0385 – gawk bug fix update

An updated gawk package that fixes three bugs is now available for Red Hat Enterprise Linux 6.

The gawk package contains the GNU version of awk, a text processing utility. AWK interprets a special-purpose programming language to do quick and easy text pattern matching and reformatting jobs.

Bug Fixes

BZ#648906

Prior to this update, the gawk utility could, under certain circumstances, interpret some run-time variables as internal zero-length variable prototypes. When gawk tried to free such run-time variables, it actually freed the internal prototypes, that were allocated just once due to memory

savings. As a consequence, gawk sometimes failed and the error message "awk: double free or corruption" was displayed. With this update the problem has been corrected and the error no longer occurs.

BZ#740673

Prior to this update, the gawk utility did not copy variables from the command line arguments. As a consequence, the variables were not accessible as intended. This update modifies the underlying code so that gawk makes copies of those variables.

BZ#743242

Prior to this update, the Yacc interpreter encountered problems handling larger stacks. As a consequence, the Yacc interpreter could fail with a stack overflow error when interpreting the AWK code. This update enlarges the stack and Yacc can now handle these AWK programs.

All users of gawk are advised to upgrade to this updated package, which fixes these bugs.

5.77. GCC

5.77.1. RHBA-2012:0941 – gcc bug fix and enhancement update

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

The gcc packages include C, C++, Java, Fortran, Objective C, Objective C++, and Ada 95 GNU compilers, along with related support libraries.

Bug Fixes

BZ#751767

The gfortran compiler could fail to compile the code with an internal compiler error. This happened because the gfc_type_for_size() function from the trans-types.c library did not return the correct data type if the demanded bit precision was less than the built-in bit precision size of the corresponding type. With this update, the function returns the corresponding wider type if no suitable narrower type has been found and the code is compiled correctly.

BZ#756138

The G++ compiler terminated unexpectedly with a segmentation fault and returned an internal compiler error when compiling with the -O2 or -O3 optimization option. This happened because the compiler tried to cancel the same loop twice in the remove_path() function. With this update, the loop is canceled only once and the segmentation fault no longer occurs in this scenario.

BZ#756651

Previously, GCC could generate incorrect code if combining instructions when splitting a two-set pattern. This was due to an error in the way the split patterns were handled while combining the instructions. With this update, the code handling instruction combining has been fixed and the problem no longer occurs.

BZ#767604

Previously, GCC could terminate unexpectedly with an internal compiler error, which was triggered by aggressive loop peeling enabled by the "-mtune=z10" setting when moving registers. With this update, the registers are determined from the instruction patterns correctly and the compilation

succeeds in this scenario.

BZ#799491

Typing into Web Console in Firefox caused Firefox to terminate unexpectedly. This happened because the compiler incorrectly cloned one of the functions called under these circumstances. With this update, the function is no longer cloned and the problem no longer occurs.

Enhancement

BZ#739443

Previously, the GCC compiler did not contain the header with functions for converting the half-float type. This update adds the header and also fixes GCC so that it works correctly with the "-march=native" option on AMD FX processor microarchitectures.

All users of gcc are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.78. GDB

5.78.1. RHBA-2012:0930 – gdb bug fix update

Updated gdb packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The GNU Debugger (GDB) is the standard debugger for Linux. With GDB, users can debug programs written in C, C++, and other languages by executing them in a controlled fashion and printing out their data.

Bug Fixes

BZ#739685

To load a core file, GDB requires the binaries that were used to produce the core file. GDB uses a built-in detection to load the matching binaries automatically. However, you can specify arbitrary binaries manually and override the detection. Previously, loading other binaries that did not match the invoked core file could cause GDB to terminate unexpectedly. With this update, the underlying code has been modified and GDB no longer crashes under these circumstances.

BZ#750341

Previously, GDB could terminate unexpectedly when loading symbols for a C++ program compiled with early GCC compilers due to errors in the cp_scan_for_anonymous_namespaces() function. With this update, an upstream patch that fixes this bug has been adopted and GDB now loads any known executables without crashing.

BZ#781571

If GDB failed to find the associated debuginfo rpm symbol files, GDB displayed the following message suggesting installation of the symbol files using the yum utility:

Missing separate debuginfo for the main executable file Try: yum --disablerepo='*' --enablerepo='*-debuginfo' install /usr/lib/debug/.build-id/47/830504b69d8312361b1ed465ba86c9e815b800

However, the suggested "--enablerepo='*-debuginfo'" option failed to work with RHN (Red Hat Network) debug repositories. This update corrects the option in the message to "--enablerepo='*-debug*'" and the suggested command works as expected.

BZ#806920

On PowerPC platforms, DWARF information created by the IBM XL Fortran compiler does not contain the DW_AT_type attribute for DW_TAG_subrange_type; however, DW_TAG_subrange_type in the DWARF information generated by GCC always contains the DW_AT_type attribute. Previously, GDB could interpret arrays from IBM XL Fortran compiler incorrectly as it was missing the DW_AT_type attribute, even though this is in accordance with the DWARF standard. This updated GDB now correctly provides a stub index type if DW_AT_type is missing for any DW_TAG_subrange_type, and processes debug info from both IBM XL Fortran and GCC compilers correctly.

All users of gdb are advised to upgrade to these updated packages, which fix these bugs.

5.79. GDM

5.79.1. RHBA-2012:1446 – gdm bug fix update

Updated gdm packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The GNOME Display Manager (GDM) is a highly configurable reimplementation of XDM, the X Display Manager. GDM allows you to log into your system with the X Window System running and supports running several different X sessions on your local machine at the same time.

Bug Fix

BZ#860646

When gdm was used to connect to a server via XDMCP (X Display Manager Control Protocol), another connection to a remote system using the "ssh -X" command resulted in wrong authorization with the X server. Consequently, applications such as xterm could not be displayed on the remote system. This update provides a compatible MIT-MAGIC-COOKIE-1 key in the described scenario, thus fixing this bug.

All users of gdm are advised to upgrade to these updated packages, which fix this bug.

5.80. GD

5.80.1. RHBA-2012:1274 – gd bug fix update

Updated gd packages that fix one bug is now available for Red Hat Enterprise Linux 6.

The gd packages provide the gd graphics library. GD allows code to draw images as PNG or JPEG files.

Bug Fix

BZ#790400

Prior to this update, ,the gd graphics library handled inverted Y coordinates incorrectly, when changing the thickness of a line. As a consequence, lines with changed thickness were drawn incorrectly. This update modifies the underlying code to draw lines with changed thickness correctly.

All users of gd are advised to upgrade to these updated packages, which fix this bug.

5.81. GEGL

5.81.1. RHSA-2012:1455 – Moderate: gegl security update

Updated gegl packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

GEGL (Generic Graphics Library) is a graph-based image processing framework.

Security Fix

CVE-2012-4433

An integer overflow flaw, leading to a heap-based buffer overflow, was found in the way the gegl utility processed .ppm (Portable Pixel Map) image files. An attacker could create a specially-crafted .ppm file that, when opened in gegl, would cause gegl to crash or, potentially, execute arbitrary code.

This issue was discovered by Murray McAllister of the Red Hat Security Response Team.

Users of gegl should upgrade to these updated packages, which contain a backported patch to correct this issue.

5.82. GERONIMO-SPECS

5.82.1. RHBA-2012:1397 – geronimo-specs bug fix update

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

The geronimo-specs packages provide the specifications for Apache's ASF-licenced J2EE server Geronimo.

Bug Fix

BZ#818755

Prior to this update, the geronimo-specs-compat package description contained inaccurate references. This update removes these references so that the description is now accurate.

All users of geronimo-specs are advised to upgrade to these updated packages, which fix this bug.

5.83. GHOSTSCRIPT

5.83.1. RHSA-2012:1256 – Moderate: ghostscript security update

Updated ghostscript packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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

Ghostscript is a set of software that provides a PostScript interpreter, a set of C procedures (the Ghostscript library, which implements the graphics capabilities in the PostScript language) and an interpreter for Portable Document Format (PDF) files.

Security Fix

CVE-2012-4405

An integer overflow flaw, leading to a heap-based buffer overflow, was found in Ghostscript's International Color Consortium Format library (icclib). An attacker could create a specially-crafted PostScript or PDF file with embedded images that would cause Ghostscript to crash or, potentially, execute arbitrary code with the privileges of the user running Ghostscript.

Red Hat would like to thank Marc Schönefeld for reporting this issue.

Users of Ghostscript are advised to upgrade to these updated packages, which contain a backported patch to correct this issue.

5.83.2. RHBA-2012:0938 – ghostscript bug fix update

Updated ghostscript packages that fix three 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#643105

Prior to this update, the gdevcups driver, which produces CUPS Raster output, handled memory allocations incorrectly. This could cause the ghostscript program to terminate unexpectedly in some situations. This update applies backported fixes for handling the memory allocations to this version of ghostscript and the crash no longer occurs.

BZ#695766

Prior to this update, certain input files containing CID Type2 fonts were rendered with incorrect character spacing. This update modifies the code so that all input files with CID Type2 fonts are rendered correctly.

BZ#697488

Prior to this update, the page orientation was incorrect when pages in the landscape orientation were converted to the PXL raster format. This update matches landscape-page sizes as well as portrait-page sizes, and sets the orientation parameter correctly when a match is found.

All users of ghostscript are advised to upgrade to these updated packages, which fix these bugs.

5.84. GIMP

5.84.1. RHSA-2012:1180 – Moderate: gimp security update

Updated gimp packages that fix three security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

The GIMP (GNU Image Manipulation Program) is an image composition and editing program.

Security Fixes

CVE-2012-3481

An integer overflow flaw, leading to a heap-based buffer overflow, was found in the GIMP's GIF image format plug-in. An attacker could create a specially-crafted GIF image file that, when opened, could cause the GIF plug-in to crash or, potentially, execute arbitrary code with the privileges of the user running the GIMP.

CVE-2011-2896

A heap-based buffer overflow flaw was found in the Lempel-Ziv-Welch (LZW) decompression algorithm implementation used by the GIMP's GIF image format plug-in. An attacker could create a specially-crafted GIF image file that, when opened, could cause the GIF plug-in to crash or, potentially, execute arbitrary code with the privileges of the user running the GIMP.

CVE-2012-3403

A heap-based buffer overflow flaw was found in the GIMP's KiSS CEL file format plug-in. An attacker could create a specially-crafted KiSS palette file that, when opened, could cause the CEL plug-in to crash or, potentially, execute arbitrary code with the privileges of the user running the GIMP.

Red Hat would like to thank Matthias Weckbecker of the SUSE Security Team for reporting the CVE-2012-3481 issue.

Users of the GIMP are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The GIMP must be restarted for the update to take effect.

5.85. GLIB2

5.85.1. RHBA-2012:0794 – glib2 bug fix update

Updated glib2 packages that fix one bug 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.

Bug Fix

BZ#782194

Prior to this upate, the gtester-report script was not marked as executable in the glib2-devel package. As a consequence, the gtester-report did not run with the default permissions. This update changes the glib2-devel package definition so that this script is now executable.

All users are advised to upgrade to these updated packages, which fix this bug.

5.86. GLIBC

5.86.1. RHBA-2012:1158 – glibc bug fix update

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

The glibc packages provide the standard C and standard math libraries used by multiple programs on the system. Without these libraries, the Linux system cannot function correctly.

Bug Fix

BZ#843571

Prior to this update, glibc incorrectly handled the "options rotate" option in the /etc/resolv.conf file when this file also contained one or more IPv6 name servers. Consequently, DNS queries could unexpectedly fail, particularly when multiple queries were issued by a single process. This update fixes internalization of the listed servers from /etc/resolv.conf into glibc's internal structures, as well as the sorting and rotation of those structures to implement the "options rotate" capability. Now, DNS names are resolved correctly in glibc in the described scenario.

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

5.86.2. RHSA-2012:1098 – Moderate: glibc security and bug fix update

Updated glibc packages that fix three security issues and one bug 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

The glibc packages provide the standard C and standard math libraries used by multiple programs on the system. Without these libraries, the Linux system cannot function properly.

Security Fix

CVE-2012-3404, CVE-2012-3405, CVE-2012-3406

Multiple errors in glibc's formatted printing functionality could allow an attacker to bypass FORTIFY_SOURCE protections and execute arbitrary code using a format string flaw in an application, even though these protections are expected to limit the impact of such flaws to an application abort.

Bug Fix

BZ#837026

A programming error caused an internal array of nameservers to be only partially initialized when the /etc/resolv.conf file contained IPv6 nameservers. Depending on the contents of a nearby structure, this could cause certain applications to terminate unexpectedly with a segmentation fault. The programming error has been fixed, which restores proper behavior with IPv6 nameservers listed in the /etc/resolv.conf file.

All users of glibc are advised to upgrade to these updated packages, which contain backported patches to fix these issues.

5.86.3. RHBA-2013:0212 – glibc bug fix update

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

The glibc packages provide the standard C and standard math libraries used by multiple programs on the system. Without these libraries, the Linux system cannot function correctly.

Bug Fix

BZ#902685

A logic error caused glibc's DNS code to incorrectly handle rejected responses from DNS servers. Consequently, after a server returned a REJECT response, additional servers defined in the /etc/resolv.conf file sometimes failed to be searched. With this update, glibc properly cycles through the servers listed in /etc/resolv.conf even if one of them returns the REJECT response, thus fixing this bug.

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

5.86.4. RHBA-2012:1422 – glibc bug fix update

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

The glibc packages provide the standard C and standard math libraries used by multiple programs on the system. Without these libraries, the Linux system cannot function correctly.

Bug Fix

BZ#864046

Prior to this update, an error in memory management within the glibc nscd daemon resulted in attempts to free a pointer that was not provided by the malloc() function. Consequently, nscd could terminate unexpectedly. This bug only happened when handling groups with a large number of members. This update ensures that memory allocated by the pool allocator is no longer passed to "free". Instead, we allow the pool allocator's garbage collector to reclaim the memory. As a result, nscd no longer crashes on groups with a large number of members.

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

5.86.5. RHSA-2012:1208 – Moderate: glibc security update

Updated glibc packages that fix multiple security issues are now available for Red Hat Enterprise Linux 6.

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

The glibc packages provide the standard C and standard math libraries used by multiple programs on the system. Without these libraries, the Linux system cannot function properly.

Security Fix

CVE-2012-3480

Multiple integer overflow flaws, leading to stack-based buffer overflows, were found in glibc's functions for converting a string to a numeric representation (strtod(), strtof(), and strtold()). If an application used such a function on attacker controlled input, it could cause the application to crash or, potentially, execute arbitrary code.

All users of glibc are advised to upgrade to these updated packages, which contain a backported patch to correct these issues.

5.86.6. RHBA-2012:0763 – glibc bug fix and enhancement update

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

The glibc packages provide the standard C and standard math libraries used by multiple programs on the system. Without these libraries, the Linux system cannot function correctly.

Bug Fixes

BZ#808545

Previously, if the **nscd** daemon received a CNAME (Canonical Name) record as a response to a DNS (Domain Name System) query, the cached DNS entry adopted the TTL (Time to Live) value of the underlying **A** or **AAAA** response. This caused the nscd daemon to wait an unexpectedly long time before reloading the DNS entry. With this update, nscd uses the shortest TTL from the response as the TTL for the entire record. DNS entries are now reloaded as expected in this scenario.

BZ#789238

Previously, locking of the main malloc arena was incorrect in the retry path. This could result in a deadlock if an sbrk request failed. With this update, locking of the main arena in the retry path has been fixed. This problem was exposed by a bug fix provided in the RHSA-2012:0058 update.

BZ#688720

glibc had incorrect information for numeric separators and groupings for French, Spanish, and German locales. Therefore, applications utilizing glibc's locale support printed numbers with the incorrect separators and groupings when those locales were in use. With this update, the separator and grouping information has been fixed.

BZ#781646

On some processors, when calling the **memcpy()** function, the optimized function variant was used. However, the optimized function variant copies the buffer backwards. As a result, if the source and target buffers were overlapping, the program behaved in an unexpected way. While such calling is a violation of ANSI/ISO standards and therefore considered an error, this update restores the prior memcpy() behavior and such programs now use the non-optimized variant of the function to allow applications to behave as before.

BZ#782585

Previously, the dynamic loader generated an incorrect ordering for initialization, which did not adhere to the ELF specification. This could result in incorrect ordering of DSO (Dynamic Shared Object) constructors and destructors. With this update, the dependency resolution has been fixed.

BZ#739971

The RHBA-2011:1179 glibc update introduced a regression, causing glibc to incorrectly parse groups with more than 126 members. Consequently, applications, such as **id**, failed to list all the groups a particular user was a member of. With this update, group parsing has been fixed.

BZ#740506

Due to a race condition within its **malloc()** routines, glibc incorrectly allocated too much memory. This could cause a multi-threaded application to allocate more memory to the threads than expected. With this update, the race condition has been fixed, and malloc's behavior is now consistent with the documentation regarding the MALLOC_ARENA_TEST and MALLOC_ARENA_MAX environment variables.

BZ#795498

Previously, **glibc** looked for an error condition in the incorrect location and therefore failed to process a second response buffer in the **gaih_getanswer()** function. As a consequence, the **getaddrinfo()** function could not properly return all addresses. This update fixes an incorrect error test condition in **gaih_getanswer()** so that **glibc** now correctly parses the second response buffer. The **getaddrinfo()** function now correctly returns all addresses.

BZ#750531

Previously, compiling code that was using the **htons()** function with the **-O2** and **-Wconversion** parameters caused bogus warnings similar to the following:

```
warning: conversion to \u2018short unsigned int\u2019 from \u2018int\u2019 may alter its value
```

This update fixes types in multiple macros and the warning is no longer returned under these circumstances.

BZ#696472

Previously, glibc did not properly detect Intel Core i3, i5, and i7 processors. As a result, glibc sometimes used incorrect implementations of several functions resulting in poor performance. This update fixes the detection process and the library provides proper function implementation to the processors.

BZ#771342

Previously, glibc did not initialize the robust futex list after a **fork()** call. As a result, shared robust mutex locks were not cleaned up after the child process exited. This update ensures that the robust futex list is correctly initialized after a fork system call.

BZ#754628

When a process corrupted its heap, the **malloc()** function could enter a deadlock while creating an error message string. As a result, the process could become unresponsive. With this update, the process uses the **mmap()** function to allocate memory for the error message instead of the **malloc()** function. The malloc() deadlock therefore no longer occurs and the process with a corrupted heap now aborts gracefully.

BZ#788959, BZ#797094, BZ#809602

Previously, glibc unconditionally used **alloca()** to allocate buffers in various routines. If such allocations applied large internal memory requests, stack overflows could occur and the application could terminate unexpectedly. This update applies several upstream patches so that glibc now uses **malloc()** for these allocations and the problem no longer occurs.

BZ#789209

Previously, glibc used an incorrect symbol for the Ukrainian currency. With this update, the symbol has been fixed.

BZ#752123

Previously, it was not possible to install the 32-bit glibc-utils package on 64-bit systems and the package was therefore missing on 64-bit Intel architectures. This update modifies the spec file so as to move the respective files and avoid conflicts. As a result, the package is now installed on these 64-bit systems as expected.

BZ#657572, BZ#785984

Previously, glibc added unneccessary spaces to abbreviated month names in the Finish and Chinese locales. With this update, the underlying code has been modified and the spaces are no longer added in the abbreviated month names in the locales.

BZ#767746

Previously, glibc returned incorrect error codes from the **pthread_create()** function. Consequently, some programs incorrectly issued an error for a transient failure, such as a temporary out-of-memory condition. This update ensures that glibc returns the correct error code when memory allocation fails in the **pthread_create()** function.

BZ#752122

Previously, glibc's dynamic loader incorrectly detected Advanced Vector Extensions (AVX) capabilities and could terminate unexpectedtly with a segmentation fault. This update fixes the AVX detection and the problem no longer occurs.

BZ#766513

Previously, an error string in glibc's **getopt** routines changed and, as the respective Japanese translation was not adapted, the system failed to find the Japanese version of the message. As a result, the error message was displayed in English even if the system locale was set to Japanese. This update fixes the Japanese translation of the error string and the problem no longer occurs.

BZ#751750

Previously, glibc's locking in the **IO_flush_all_lockp()** function was incorrect. This resulted in a race condition with occasional deadlocks when calling the **fork()** function in multi-threaded applications. This update fixes the locking and avoids the race condition.

BZ#784402

Previously, the **nscd** daemon cached all transient results even if they were negative. This could result in erroneous nscd results. This update ensures that negative results of transient errors are not cached.

BZ#804630

When the **resolv.conf** file contained only nameservers with IPv6 and **options rotate** was set, the search domain was always appended. However, this is not desired in the case of fully qualified domain names (FQDN) and if an FQDN was used, the resolution failed. With this update, the underlying code has been modified and if more than one IPv6 nameserver is defined in **resolv.conf**, the FQDN is resolved correctly. Refer to bug 771204 for further information about this problem.

BZ#789189

Previously, when parsing the **resolv.conf** file, glibc did not handle the parsing of spaces in nameserver entries correctly. Consequently, correct DNS lookups failed. This update fixes the space parsing and the problem no longer occurs.

BZ#804689

The **getaddrinfo()** call could return an incorrect value. This happened because the query for getaddrinfo was more complex than necessary and getaddrinfo failed to handle the additional information returned by the query correctly. With this update, the query no longer returns the addition information and the problem is fixed.

Enhancements

BZ#697421, BZ#749188

Previously, glibc did not support the ISO-10646-UCS-2 character set for the following locales: az_AZ, as_IN, and tt_RU. This update adds support for the character set and the locales.

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

5.87. GNOME-DESKTOP

5.87.1. RHBA-2012:1352 – gnome-desktop bug fix update

Updated gnome-desktop packages that fix a bug are now available.

The gnome-desktop package contains an internal library (libgnome-desktop) used to implement some portions of the GNOME desktop, and also some data files and other shared components of the GNOME user environment.

Bug Fix

BZ#829891

Previously, when a user hit the system's hot-key (most commonly Fn+F7) to change display configurations, the system could potentially switch to an invalid mode, which would fail to display. With this update, gnome-desktop now selects valid XRandR modes and correctly switching displays with the hot-key works as expected.

All users of gnome-desktop are advised to upgrade to these updated packages, which fix this bug.

5.87.2. RHBA-2012:0405 – gnome-desktop bug fix update

An updated gnome-desktop package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The gnome-desktop package contains an internal library (libgnomedesktop) used to implement some portions of the GNOME desktop, and also some data files and other shared components of the GNOME user environment.

Bug Fix

BZ#639732

Previously, due to an object not being destroyed, the Nautilus file manager could consume an excessive amount of memory. Consequently, constantly growing resident memory would slow down the system. The source code has been modified to prevent memory leaks from occurring and Nautilus now consumes a reasonable amount of memory.

All users of gnome-desktop are advised to upgrade to this updated package, which fixes this bug.

5.88. GNOME-KEYRING

5.88.1. RHBA-2012:1334 – gnome-keyring bug fix update

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

The gnome-keyring session daemon manages passwords and other types of secrets for the user, storing them encrypted with a main password. Applications can use the gnome-keyring library to integrate with the key ring.

Bug Fix

BZ#860644

Due to a bug in the thread-locking mechanism, the gnome-keyring daemon could sporadically become unresponsive while reading data. This update fixes the thread-locking mechanism and no more deadlocks occur in gnome-keyring in the described scenario.

All gnome-keyring users are advised to upgrade to these updated packages, which fix this bug.

5.88.2. RHBA-2012:0878 – gnome-keyring bug fix update

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

The gnome-keyring session daemon manages passwords and other types of secrets for the user, storing them encrypted with a main password. Applications can use the gnome-keyring library to integrate with the keyring.

Bug Fixes

BZ#708919, BZ#745695

Previously, the mechanism for locking threads was missing. Due to this, gnome-keyring could have, under certain circumstances, terminated unexpectedly on multiple key requests from the integrated ssh-agent. With this update, the missing mechanism has been integrated into gnome-keyring so that gnome-keyring now works as expected.

All users of gnome-keyring are advised to upgrade to these updated packages, which fix these bugs.

5.89. GNOME-PACKAGEKIT

5.89.1. RHBA-2012:1229 – gnome-packagekit bug fix update

Updated gnome-packagekit packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The gnome-packagekit packages provide session applications for the PackageKit API.

Bug Fix

BZ#839197

Previously, it was possible for the user to log out of the system or shut it down while the PackageKit update tool was running and writing to the RPM database (rpmdb). Consequently, rpmdb could become damaged and inconsistent due to the unexpected termination and cause various problems with subsequent operation of the rpm, yum, and PackageKit utilities. This update modifies PackageKit to not allow shutting down the system when a transaction writing to rpmdb is active, thus fixing this bug.

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

5.90. GNOME-POWER-MANAGER

5.90.1. RHBA-2012:0935 – gnome-power-manager bug fix update

Updated gnome-power-manager packages that fix one bug are now available for Red Hat Enterprise Linux 6.

GNOME Power Manager uses the information and facilities provided by UPower displaying icons and handling user callbacks in an interactive GNOME session.

Bug Fix

BZ#676866

After resuming the system or re-enabling the display, an icon could appear in the notification area with an erroneous tooltip that read "Session active, not inhibited, screen idle. If you see this test, your display server is broken and you should notify your distributor." and included a URL to an external web page. This error message was incorrect, had no effect on the system and could be safely ignored. In addition, linking to an external URL from the notification and status area is unwanted. To prevent this, the icon is no longer used for debugging idle problems.

All users of gnome-power-manager are advised to upgrade to these updated packages, which fix this bug.

5.91. GNOME-SCREENSAVER

5.91.1. RHBA-2012:1393 – gnome-screensaver bug fix update

Updated gnome-screensaver packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The gnome-screensaver packages contain the GNOME project's official screen saver program. It 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 Fix

BZ#860643

When a Mandatory profile was enabled, the "Lock screen when screen saver is active" option in the Gnome Screensaver Preferences window was not disabled. This bug could lead to security risks for users. With this update, the lock-screen option is disabled as expected in the described scenario, thus preventing this bug.

All users of gnome-screensaver are advised to upgrade to these updated packages, which fix this bug.

5.92. GNOME-SETTINGS-DAEMON

5.92.1. RHBA-2012:1450 – gnome-settings-daemon bug fix update

Updated gnome-settings-daemon packages that fix a bug is now available for Red Hat Enterprise Linux 6.

The gnome-settings-daemon packages contain a daemon to share settings from GNOME with other applications. It also handles global key bindings, as well as a number of desktop-wide settings.

Bug Fix

BZ#866528

Previously, when a system hotkey was used to change the display configuration, sometimes a valid XRandR configuration failed to be selected and the monitors were not kept in clone mode. Consequently, it was impossible to switch displays. With this update, gnome-settings-daemon always selects valid XRandR modes, and sets or unsets clone mode as expected, thus fixing this bug.

Users of gnome-settings-daemon are advised to upgrade to these updated packages, which fix this bug.

5.92.2. RHBA-2012:0949 – gnome-settings-daemon bug fix and enhancement update

An updated gnome-settings-daemon package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The gnome-settings-daemon package contains a daemon to share settings from GNOME with other applications. It also handles global key bindings, as well as a number of desktop-wide settings.

Bug Fixes

BZ#693843

Previously, the selected keyboard layout on certain machines reverted to the "US" layout every time the user logged in. With this update, the bug has been fixed so that the selected keyboard layout is not reverted anymore.

BZ#805036

Previously, the automatic mapping of the screen tablet did not work with the NVIDIA driver. With this update, support for the NV-CONTROL extension has been added so that the automatic mapping of the screen tablet now works as expected.

BZ#805042

Previously, the button mapping to actions did not work in the Wacom graphics tablet plug-in. As a result, the Map Buttons did not display in the GUI and activating buttons on the Wacom graphics tablet had no effect. With this update, these problems have been fixed.

Enhancements

BZ#769464

With this update, Wacom graphics tablets are now supported with gnome-settings-daemon.

BZ#816646

This update modifies the way gnome-settings-daemon stores settings in GConf. Previously, the settings were stored per user and per device. With this update, the settings are now stored per user, per device, and per machine.

All users of gnome-settings-daemon are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.93. GNOME-SYSTEM-MONITOR

5.93.1. RHBA-2012:0769 – gnome-system-monitor bug fix update

An updated gnome-system-monitor package that fixes two bugs is now available for Red Hat Enterprise Linux 6.

The gnome-system-monitor utility allows users to graphically view and manipulate the running processes on the system, and provides an overview of available resources such as CPU and memory.

Bug Fixes

BZ#682011

Prior to this update, the gnome-system-monitor failed to correctly parse the contents of the /proc/cpuinfo file if it included an informational entry about the machine model on 64-bit PowerPC architectures. As a consequence, a false "Unknown CPU model" processor was incorrectly reported by the application. This update changes the parsing code to discard such information when it does not identify an additional processor.

BZ#692956

Prior to this update, the gnome-system-monitor parser code expected a certain string to identify the CPU speed which is not used for all architectures. As a consequence, the gnome-system-monitor could fail to correctly parse the processor speed from /proc/cpuinfo when a different string was used, for example on 64-bit PowerPC. This update changes the parsing code to support different string types used on such architectures.

All users of gnome-system-monitor are advised to upgrade to this updated package, which fixes these bugs.

5.94. GNOME-TERMINAL

5.94.1. RHBA-2012:1311 – gnome-terminal bug fix update

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

Gnome-terminal is a terminal emulator for GNOME. It supports translucent backgrounds, opening multiple terminals in a single window (tabs) and clickable URLs.

Bug Fix

BZ#819796

Prior to this update, gnome-terminal was not completely localized into Asamese. With this update, the Assamese locale has been updated.

All gnome-terminal users are advised to upgrade to these updated packages, which fix this bug.

5.95. GRAPHVIZ

5.95.1. RHBA-2012:1291 – graphviz bug fix update

Updated graphviz packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

Graphviz is open-source graph-visualization software. Graph visualization is a way of representing structural information as diagrams of abstract graphs and networks. It has important applications in networking, bioinformatics, software engineering, database and web design, machine learning, and in visual interfaces for other technical domains.

Bug Fixes

BZ#772637

Previously, the dot tool could generate different images on 32-bit and 64-bit architectures, which could consequently lead to multilib conflicts of packages that use graphviz during its build process. The problem was caused by different instructions used for floating points processing. On 32-bit Intel architecture, the code is now compiled with the "--ffloat-store" compiler flag, which ensures that identical images are generated regardless of the used architecture.

BZ#821920

The graphviz-tcl package included the "demo" directory, which contained examples in various languages. This caused implicit dependencies to be introduced. With this update, all examples are installed as documentation, which reduces the number of implicit dependencies.

BZ#849134

The "dot -c" command which is run in the %postun scriptlet recreates graphviz configuration files to be up-to-date with the current state of the installed plug-ins. Previously, if the command failed to load plug-ins specified in the configuration files, warning messages were printed when removing the graphviz-gd package. These messages could have been confusing, and have been therefore removed.

All users of graphviz are advised to upgrade to these updated packages, which fix these bugs.

5.96. GREP

5.96.1. RHBA-2012:0352 – grep bug fix update

An updated grep package that fixes one bug is 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 Fix

BZ#741452

Previously, the grep utility was not able to handle the EPIPE error. If a SIGPIPE signal was blocked by the shell, grep kept continuously printing error messages. An upstream patch has been applied to address this problem, so that grep exits on the first EPIPE error and prints only one error message.

All users of grep are advised to upgrade to this updated package, which fixes this bug.

5.97. GRUBBY

5.97.1. RHBA-2012:0895 – grubby bug fix and enhancement update

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

The grubby packages provide grubby, a command line tool for displaying and editing of GRUB (GRand Unified Bootloader) configuration files.

Bug Fix

BZ#696960

Previously, when grubby was executed with the "--args=[arguments] --update-kernel=ALL" options to update command line arguments for all kernels whose boot configuration was stored in the edited configuration file, it updated only arguments for the first kernel in the file. As a result, arguments for the other kernels were not updated. This update ensures that arguments for all kernels in a configuration file are updated when grubby is launched with the aforementioned options.

All users of grubby are advised to upgrade to these updated packages, which fix this bug.

5.98. GRUB

5.98.1. RHBA-2012:0892 – grub bug fix update

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

The GRUB utility is responsible for booting the operating system kernel.

Bug Fix

BZ#670266

Due to an error in the underlying source code, previous versions of GRUB sometimes failed to boot in Unified Extensible Firmware Interface (UEFI) mode when booting from the network on systems with multiple Pre-boot Execution Environment (PXE) network interface cards (NICs). This update ensures that GRUB attempts to identify and use an active interface that has already successfully acquired an address via Dynamic Host Configuration Protocol (DHCP) instead of using the one suggested by the system. As a result, booting from the network in UEFI mode now works as expected on systems with multiple NICs.

All users of GRUB are advised to upgrade to these updated packages, which fix this bug.

5.99. GSTREAMER-PLUGINS-BASE

5.99.1. RHEA-2012:1473 — gstreamer-plugins-base enhancement update

Updated gstreamer-plugins-base packages thatadd one enhancement are now available for Red Hat Enterprise Linux 6.

The gstreamer-plugins-base packages provide a collection of base plug-ins for the GStreamer streaming media framework.

Enhancement

BZ#755777

This update adds color-matrix support for color conversions to the ffmpegcolorspace plugin.

All users of gstreamer-plugins-base are advised to upgrade to these updated packages, which add this enhancement.

5.100. GTK2

5.100.1. RHBA-2012:0809 — gtk2 bug fix and enhancement update

Updated gtk2 packages that fix three bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

GTK+ is a multi-platform toolkit for creating graphical user interfaces.

Bug Fixes

BZ#697437

Previously, the "Open Files" dialog box failed to show the "Size" column if it was previously used in "Search" mode. This update fixes the bug by ensuring that the "Size" column is always displayed accordingly to the "Show Size Column" context menu option.

BZ#750756

Previously, copying text from selectable labels, such as those displayed in message dialog boxes, using the Ctrl+Insert key combination did not work. This update adds the Ctrl+Insert key combination that copies selected text to clipboard when activated.

BZ#801620

Previously, certain GTK applications, such as virt-viewer, failed to properly initialize key bindings associated with menu items. This was due to a bug in the way properties associated with the menu items were parsed by the library. This update fixes the bug, rendering the menu items accessible again by key bindings for applications that use this feature.

Enhancement

BZ#689188

Previously, the "Open Files" dialog box could appear with an abnormal width when the "file type" filter contained a very long string (as observed with certain image hosting websites), making the dialog unusable. With this update, the dialog box splits the filter string into multiple lines of text, so that the dialog keeps a reasonable width.

All users of gtk2 are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.101. GVFS

5.101.1. RHBA-2012:1124 – gvfs bug fix and enhancement update

Updated gvfs packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

GVFS is the GNOME desktop's virtual file system layer, which allows users to easily access local and remote data, including via the FTP, SFTP, WebDAV, CIFS and SMB protocols, among others. GVFS integrates with the GIO (GNOME I/O) abstraction layer.

Bug Fixes

BZ#599055

Previously, rules for ignoring mounts were too restrictive. If the user clicked on an encrypted volume in the Nautilus' sidebar, an error message was displayed and the volume could not be accessed. The underlying source code now contains additional checks so that encrypted volumes have proper mounts associated (if available), and the file system can be browsed as expected.

BZ#669526

Due to a bug in the kernel, a freshly formatted Blu-ray Disk Rewritable (BD-RE) medium contains a single track with invalid data that covers the whole medium. This empty track was previously incorrectly detected, causing the drive to be unusable for certain applications, such as Brasero. This update adds a workaround to detect the empty track, so that freshly formatted BD-RE media are properly recognized as blank.

BZ#682799, BZ#746977, BZ#746978, BZ#749369, BZ#749371, BZ#749372

The code of the gvfs-info, gvfs-open, gvfs-cat, gvfs-ls and gvfs-mount utilities contained hard-coded exit codes. This caused the utilities to always return zero on exit. The exit codes have been revised so that the mentioned gvfs utilities now return proper exit codes.

BZ#746905

When running gvfs-set-attribute with an invalid command-line argument specified, the utility terminated unexpectedly with a segmentation fault. The underlying source code has been modified so that the utility now prints a proper error message when an invalid argument is specified.

BZ#809708

Due to missing object cleanup calls, the gvfsd daemon could use excessive amount of memory, which caused the system to become unresponsive. Proper object cleanup calls have been added with this update, which ensures that the memory consumption is constant and the system does not hang in this scenario.

All users of gvfs are advised to upgrade to these updated packages, which fix these bugs.

5.102. HIVEX

5.102.1. RHBA-2012:0776 – hivex bug fix and enhancement update

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

Hive files are undocumented binary files that Windows uses to store the Windows Registry on the disk. Hivex is a library that can read and write to these files.

The hivex packages have been upgraded to upstream version 1.3.3, which provides a number of bug fixes and enhancements over the previous version. (BZ#734208)

All hivex users are advised to upgrade to these updated hivex packages, which fix these bugs and add these enhancements.

5.103. HSQLDB

5.103.1. RHBA-2012:0993 – hsqldb enhancement update

Updated hsqldb packages that add an enhancement are now available for Red Hat Enterprise Linux 6.

HSQLDB is a relational database engine written in Java, with a JDBC driver, supporting a subset of ANSI-92 SQL. It offers a small (about 100k), fast database engine which offers both in-memory and disk-based tables. Embedded and server modes are available. Additionally, it includes tools such as a minimal web server, in-memory query and management tools (which can be run as applets or servlets), and a number of demonstration examples.

Enhancement

BZ#816735

HSQLdb has been updated to add stubs for JDBC 4.1

Users of hsqldb are advised to upgrade to these updated packages, which add this enhancement.

5.104. HWDATA

5.104.1. RHEA-2012:0879 – hwdata enhancement update

An updated hwdata package that adds two enhancements is now available for Red Hat Enterprise Linux 6.

The hwdata package contains tools for accessing and displaying hardware identification and configuration data.

Enhancements

BZ#737467

With this update, the monitor database has been updated with information about the Acer 76ie monitor. Also, several duplicate monitor entries have been removed from the database.

BZ#760014

The pci.ids database has been updated with information about the Atheros wireless network adapter, Killer Wireless-N 1103.

All users of hwdata are advised to upgrade to this updated package, which adds these enhancements.

5.105. ICEDTEA-WEB

5.105.1. RHSA-2012:1132 – Important: icedtea-web security update

Updated icedtea-web packages that fix two security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

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.

Security Fixes

CVE-2012-3422

An uninitialized pointer use flaw was found in the IcedTea-Web plug-in. Visiting a malicious web page could possibly cause a web browser using the IcedTea-Web plug-in to crash, disclose a portion of its memory, or execute arbitrary code.

CVE-2012-3423

It was discovered that the IcedTea-Web plug-in incorrectly assumed all strings received from the browser were NUL terminated. When using the plug-in with a web browser that does not NUL terminate strings, visiting a web page containing a Java applet could possibly cause the browser to crash, disclose a portion of its memory, or execute arbitrary code.

Red Hat would like to thank Chamal De Silva for reporting the CVE-2012-3422 issue.

This erratum also upgrades IcedTea-Web to version 1.2.1. Refer to the NEWS file for further information.

All IcedTea-Web users should upgrade to these updated packages, which resolve these issues. Web browsers using the IcedTea-Web browser plug-in must be restarted for this update to take effect.

5.105.2. RHSA-2012:1434 – Critical: icedtea-web security update

Updated icedtea-web packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

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.

Security Fix

CVE-2012-4540

A buffer overflow flaw was found in the IcedTea-Web plug-in. Visiting a malicious web page could cause a web browser using the IcedTea-Web plug-in to crash or, possibly, execute arbitrary code.

Red Hat would like to thank Arthur Gerkis for reporting this issue.

This erratum also upgrades IcedTea-Web to version 1.2.2. Refer to the NEWS file for further information:

http://icedtea.classpath.org/hg/release/icedtea-web-1.2/file/icedtea-web-1.2.2/NEWS

All IcedTea-Web users should upgrade to these updated packages, which resolve this issue. Web browsers using the IcedTea-Web browser plug-in must be restarted for this update to take effect.

5.105.3. RHBA-2012:0845 – 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.

IcedTea-Web provides a Java web browser plug-in, a Java Web Start implementation, and the IcedTea Web Control Panel.

The icedtea-web packages have been upgraded to upstream version 1.2, which provides a number of bug fixes and enhancements over the previous version. (BZ#756843)

Note: This update is not compatible with Firefox 3.6 and earlier. If you are using such a Firefox version, upgrade to a later supported version before applying this update.

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

5.106. IMSETTINGS

5.106.1. RHBA-2012:0768 – imsettings bug fix update

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

IMSettings provides command line tools and a library to configure and control input-methods settings. Users normally access it through the "im-chooser" GUI tool.

Bug Fix

BZ#713433

Prior to this update, the IMSettings daemon unexpectedly invalidated the previous pointer after obtaining a new pointer. This update modifies IMSettings so that the code is updated after all transactions are finished.

All users of imsettings are advised to upgrade to these updated packages, which fix this bug.

5.107. INDENT

5.107.1. RHBA-2012:0753 – indent bug fix update

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

The indent package provides a utility to convert from one C writing style to a different one. Indent understands correct C syntax and can handle incorrect C syntax.

Bug Fixes

BZ#733265

Prior to this update, suffixes were incorrectly separated when running the indent utility on code with decimal float constants. As a consequence, indent could encounter a compilation syntax error. This update modifies indent to understand decimal float suffixes as proposed by the N1312 draft of ISO/IEC WDTR24732. Now, indent handles decimal float constants as expected.

BZ#784304

Prior to this update, the internal test-suite did not signal test failure by exit code if indent failed to pass the test. This update adds an exit call with non-zero value to signal failure.

All users of indent are advised to upgrade to this updated package, which fixes these bugs.

5.108. INITSCRIPTS

5.108.1. RHBA-2012:1275 – initscripts bug fix update

Updated initscripts packages that fix one bug 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 Fix

BZ#854852

Previously, the naming policy for VLAN names was too strict. Consequently, the if-down utility did not properly remove descriptively-named interfaces from the /proc/net/vlan/config file. This update removes the name format check and if-down now works as expected in the described scenario.

All users of initscripts are advised to upgrade to these updated packages, which fix this bug.

5.108.2. RHBA-2012:0816 – initscripts bug fix and enhancement update

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

The initscripts package contains system scripts to boot your system, change runlevels, activate and deactivate most network interfaces, and shut the system down cleanly.

Bug Fixes

BZ#781493

The previous version of initscripts did not support IPv6 routing in the same way as IPv4 routing. IPv6

addressing and routing could be achieved only by specifying the **ip** commands explicitly with the **-6** flag in the **/etc/sysconfig/network-scripts/rule-*DEVICE_NAME*** configuration file where *DEVICE_NAME* is the name of the respective network interface. With this update, the related network scripts have been modified to provide support for IPv6-based policy routing and IPv6 routing is now configured separately in the **/etc/sysconfig/network-scripts/rule6-*DEVICE_NAME*** configuration file.

BZ#786404

During the first boot after system installation, the kernel entropy was relatively low to generate high-quality keys for **sshd**. With this update, the entropy created by the disk activity during system installation is saved in the **/var/lib/random-seed** file and used for key generation. This provides enough randomness and allows generation of keys based on sufficient entropy.

BZ#582002

In emergency mode, every read request from the **/dev/tty** device ended with an error and consequently, it was not possible to read from the **/dev/tty** device. This happened because, when activating single-user mode, the **rc.sysinit** script called the **sulogin** application directly. However, sulogin needs to be the console owner to operate correctly. With this update, rc.sysinit starts the **rcS-emergency** job, which then runs **sulogin** with the correct console setting.

BZ#588993

The **ifconfig** utility was not able to handle 20-byte MAC addresses in InfiniBand environments and reported that the provided addresses were too long. With this update, the respective **ifconfig** commands have been changed to aliases to the respective **ip** commands and **ifconfig** now handles 20-byte MAC addresses correctly.

BZ#746045

Due to a logic error, the **sysfs()** call did not remove the ***arp_ip_target*** correctly. As a consequence, the following error was reported when attempting to shut down a bonding device:

```
ifdown-eth: line 64: echo: write error: Invalid argument
```

This update modifies the script so that the error no longer occurs and ***arp_ip_target*** is now removed correctly.

BZ#746808

The **serial.conf** file now contains improved comments on how to create an **/etc/init/tty<device>.conf** file that corresponds to the active serial device.

BZ#802119

The **network** service showed error messages on service startup similar to the following:

```
Error: either "dev" is duplicate, or "20" is a garbage.
```

This was due to incorrect splitting of the parsed arguments. With this update, the arguments are processed correctly and the problem no longer occurs.

BZ#754984

The **halt** initscript did not contain support for the **apcupsd** daemon, the daemon for power mangement and controlling of APC's UPS (Uninterruptible Power Supply) supplies. Consequently, the supplies were not turned off on power failure. This update adds the support to the script and the

UPS models are now turned off in power-failure situations as expected.

BZ#755175

In the previous version of initscripts, the comments with descriptions of variables **kernel.msgmnb** and **kernel.msgmax** were incorrect. With this update, the comments have been fixed and the variables are now described correctly.

BZ#787107

Due to an incorrect logic operator, the following error was returned on network service shutdown as the shutdown process failed:

```
69: echo: write error: Invalid argument
```

With this update, the code of the shutdown initscript has been modified and the error is no longer returned on network service shutdown.

BZ#760018

The system could remain unresponsive for some time during shutdown. This happened because initscript did not check if there were any CIFS (Common Internet File System) share mounts and failed to unmount any mounted CIFS shares before shutdown. With this update, a CIFS shares check has been added and the shares are stopped prior to shutdown.

BZ#721010

The **ifup-aliases** script was using the **ifconfig** tool when starting IP alias devices. Consequently, the **ifup** execution was gradually slowing down significatly with the increasing number of the devices on the NIC (Network Interface Card) device. With this update, IP aliases now use the **ip** tool instead of **ifconfig** and the performance of the **ifup-aliases** script remains constant in the scenario described.

BZ#765835

Prior to this update, the **netconsole** script could not discover and resolve the MAC address of a router specified in the **/etc/sysconfig/netconsole** file. This happened because the address was resolved as two identical addresses and the script failed. This update modifies the **netconsole** script so that it handles the MAC address correctly and the device is discovered as expected.

BZ#757637

In the Malay (**ms_MY**) locale, some services did not work properly. This happened due to a typographical mistake in the ms.po file. This update fixes the mistake and services in the ms_MY locale run as expected.

BZ#749610

The **primary** option for bonding in the **ifup-eth** tool had a timing issue when bonding NIC devices. Consequently, the bonding was configured, but it was the active interface that was enslaved first. With this update, the timing of bonding with the **primary** option has been corrected and the device defined in the **primary** option is enslaved first as expected.

Enhancement

BZ#704919

Users can now set the NIS (Network Information Service) domain name by configuring the ***NISDOMAIN*** parameter in the /etc/sysconfig/network file, or other relevant configuration files.

Users of initscripts should upgrade to these updated packages, which fix these bugs and add this enhancement.

5.109. IOK

5.109.1. RHBA-2012:1164 – iok bug fix update

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

The iok package contains an Indic on-screen virtual keyboard that supports the Assamese, Bengali, Gujarati, Hindi, Kannada, Marathi, Malayalam, Punjabi, Oriya, Sindhi, Tamil and Telugu languages. Currently, iok works with Inscript and xkb keymaps for Indian languages, and is able to parse and display non-Inscript keymaps as well.

Bug Fixes

BZ#814541, BZ#814548

Previously, when saving a keymap with a specified name, predefined naming convention was followed and the file name was saved with the "-" prefix without noticing the user. With this update, if the user attempts to save a keymap, a dialog box displaying the required file name format appears.

BZ#819795

This update provides the complete iok translation for all supported locales.

All users of iok are advised to upgrade to these updated packages, which fix these bugs.

5.109.2. RHBA-2012:0392 – iok bug fix update

An updated iok package that fixes multiple bugs is now available for Red Hat Enterprise Linux 6.

The iok package contains an Indic on-screen virtual keyboard that supports the Assamese, Bengali, Gujarati, Hindi, Kannada, Marathi, Malayalam, Punjabi, Oriya, Sindhi, Tamil and Telugu languages. Currently, iok works with Inscript and xkb keymaps for Indian languages, and is able to parse and display non-Inscript keymaps as well.

The iok package has been upgraded to upstream version 1.3.13, which provides a number of bug fixes over the previous version.

Bug Fixes

BZ#736992

Due to xkb keymaps being rewritten in a recent update of the xkeyboard-config package, the iok's language list contained incorrect xkb keymap names when selecting the Hindi X Keyboard Extension (XKB). To fix this problem, the iok's xkb parser has been rewritten.

BZ#752667

Previously, iok looked for files with the ".mim" suffix in the "~/.m17n" directory instead of the "~/.m17n.d" directory. This update modifies the directory path to the correct "~/.m17n.d" so that the user-defined keymap files are saved in the correct directory.

BZ#752668

Previously, when using the on-screen keyboard, mouse clicks on various characters worked as expected. However, finger inputs failed because the first selected character was selected regardless of what characters the user selected next. With this update, users can use the drag-and-drop feature when running iok in advanced mode (the "iok -a" command), which allows users to drag the first key button over the second button. The drag-and-drop feature is not available in iok's default mode.

BZ#798592

Due to a small size of the xkb name array, if the user selected the xkb-Malayalam keymap (enhanced Indian Script with the Rupee sign), and then pressed the "To English" button, the iok utility could terminate unexpectedly. With this update, the size of the xkb name array has been increased so that the utility no longer crashes in the described scenario.

All users of iok are advised to upgrade to this updated package, which fixes these bugs.

5.110. IPA

5.110.1. RHSA-2013:0188 – Important: ipa security update

Updated ipa packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

Red Hat Identity Management is a centralized authentication, identity management and authorization solution for both traditional and cloud-based enterprise environments.

Security Fix

CVE-2012-5484

A weakness was found in the way IPA clients communicated with IPA servers when initially attempting to join IPA domains. As there was no secure way to provide the IPA server's Certificate Authority (CA) certificate to the client during a join, the IPA client enrollment process was susceptible to man-in-the-middle attacks. This flaw could allow an attacker to obtain access to the IPA server using the credentials provided by an IPA client, including administrative access to the entire domain if the join was performed using an administrator's credentials.



NOTE

This weakness was only exposed during the initial client join to the realm, because the IPA client did not yet have the CA certificate of the server. Once an IPA client has joined the realm and has obtained the CA certificate of the IPA server, all further communication is secure. If a client were using the OTP (one-time password) method to join to the realm, an attacker could only obtain unprivileged access to the server (enough to only join the realm).

Red Hat would like to thank Petr Menšík for reporting this issue.

This update must be installed on both the IPA client and IPA server. When this update has been applied to the client but not the server, ipa-client-install, in unattended mode, will fail if you do not have the correct CA certificate locally, noting that you must use the "--force" option to insecurely obtain the

certificate. In interactive mode, the certificate will try to be obtained securely from LDAP. If this fails, you will be prompted to insecurely download the certificate via HTTP. In the same situation when using OTP, LDAP will not be queried and you will be prompted to insecurely download the certificate via HTTP.

Users of ipa are advised to upgrade to these updated packages, which correct this issue. After installing the update, changes in LDAP are handled by ipa-ldap-updater automatically and are effective immediately.

5.110.2. RHBA-2012:0819 – ipa bug fix and enhancement update

Updated ipa packages that fix several bugs and add various 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.



NOTE

The ipa package has been upgraded to upstream version 2.2.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#736865)

Bug Fixes

BZ#810900

The Identity Management password policy plug-in for the Directory Server did not properly sort the history of user passwords when it was checking the sanity of a password change. Due to this bug, the user password history was sorted randomly, and, consequently, a random password was removed rather than the oldest password when the list overflowed. As a result, users could bypass the password policy requirement for password repetition. User passwords are now sorted correctly in the Identity Management password plug-in for the Directory Server, and the password policy requirement for password repetition is properly enforced.

BZ#805478

Due to a bug in the Identity Management permission plug-in, an attempt to rename a permission always resulted in an error. Consequently, users had to remove the permission and create a new permission with a new name when attempting to rename a permission. With this update, the underlying source code has been modified to address this issue, and users are now able to rename permissions.

BZ#701677

Previously, the DNS plug-in did not allow users to set a query or a transfer policy for a zone managed by Identity Management. Therefore, users could not control who could query or transfer zones in the same way they do with zones stored in plain text files. With this update, users can set ACLs for every zone managed by Identity Management; thus, users can control who can query their zones or run zone transfers.

BZ#773759

Non-admin users with an appropriate permission can change passwords of other users. However, the target group of this permission was previously not limited. Consequently, a non-admin user with the

permission to change passwords could change the password of the admin user and acquire access to the admin account. With this update, the permission was changed to allow password changes for non-admin users only.

BZ#751173

When the **ipa passwd** CLI command was used to change user's password, it returned the following error message when the password change failed:

> ipa: ERROR: Constraint violation: Password Fails to meet minimum strength criteria

User password changes are a subject of a configured password policy. Without a proper error message, it may be difficult to investigate why the password change failed (password complexity, too soon to change password, etc.) and amend the situation. The Directory Server plug-in that is used to change passwords now returns a proper error message if the **ipa passwd** command fails.

BZ#751597

When an Identity Management server is installed with a custom hostname which is not properly resolvable in DNS, an IP address for the custom hostname is requested from the user. Next, a host record is added to the **/etc/hosts** file so that the custom hostname is resolvable and the installation can continue. However, previously, the record was not added when the IP address was passed using the **--ip-address** option. As a result, installation failed because subsequent steps could not resolve the machine's IP address. With this update, a host record is added to **/etc/hosts** even when the IP address is passed via the **--ip-address** option, and the installation process continues as expected.

BZ#751769

Identity Management could not be installed on a server with a custom LDAP server instance even though the LDAP server instance runs on a custom port and therefore does not conflict with Identity Management. As a result, users could not deploy custom LDAP instances on a system with Identity Management. With this update, Identity Management no longer enforces that no LDAP instances exist. Instead, it checks that reserved LDAP ports (**389** and **636**) are free. Users can combine an Identity Management server with custom LDAP server instances as long as they run on custom ports.

BZ#753484

When the Kerberos single sign-on to the Identity Management Web UI failed, the Web UI did not fall back to the login and password authentication. Workstations outside of the Identity Management Kerberos realm, or with incompatible browsers, could not access the Web UI unless a fallback from Kerberos authentication to login and password authentication was configured on the Identity Management web server. The Web UI is now able to fall back to form based authentication when Kerberos authentication cannot be used.

BZ#754973

The **force-sync**, **re-initialize**, and **del** sub-commands of the **ipa-replica-manage** command failed when used against a winsync agreement on an Active Directory machine, limiting the user's ability to control winsync replication agreements. With this update, the **ipa-replica-manage** was fixed to manage both standard replication agreement and winsync agreements in a more robust way.

BZ#757681

The Identity Management installer did not process the host IP address properly when the **--no-host-dns** option was passed. When a hostname was not resolvable and the **--no-host-dns** option was used, the **ipa-replica-install** utility failed during the installation and did not amend the hostname

resolution in the same way as the **ipa-server-install** utility does. With this update, **ipa-server-install** and **ipa-replica-install** now share host IP address processing, and both add a record to the **/etc/hosts** file when the server or replica hostname is not resolvable.

BZ#759100

The Identity Management server installation script did not properly handle situations when a server had 2 IP addresses assigned, and failed to proceed with the installation. This update fixes the installation script, and installing the Identity Management server in a dual-NIC configuration works as expected.

BZ#750828

When Identity Management is installed with the **--external-ca** option, the installation is divided in two stages. The second stage of the installation process reads configuration options from a file stored by the first stage. Previously, the installer did not properly store a value with the DNS forwarder IP address, which was then misread by the second stage of the installation process, and name server configuration in the second stage of the installation failed. With this update, the forwarder option is correctly stored, and installation works as expected.

BZ#772043

Prior to this update, the Identity Management netgroup plug-in did not validate netgroup names. Consequently, a netgroup with an invalid name could be stored in an LDAP server which could then crash when the invalid value was processed by the NIS plug-in. The Identity Management **netgroup** plug-in now enforces stricter validation of netgroup names.

BZ#772150

Certain Identity Management replica agreements ignored a list of attributes that should have been excluded from replication. Identity Management attributes that are generated locally on each master by the LDAP server plug-in (in this case, the **memberOf** attribute) were being replicated. This forced all Identity Management replicas' LDAP servers to re-process the **memberOf** data and increase the load on the LDAP servers. When many entries were added to a replica in a short period of time, or when a replica was being re-initialized from another master, all replicas were flooded with **memberOf** changes, which caused high load on all replica machines and caused performance issues. New replica agreements, added by the **ipa-replica-install** utility, no longer ignore lists of attributes excluded from replication. Re-initialization or a high number of added entries in an Identity Management LDAP server no longer causes performance issues caused by **memberOf** processing. Old replica agreements have also been updated to contain the correct list of attributes excluded from replication.

BZ#784025

The **ipa automountmap-add-indirect** command creates a new map and adds a key to the parent map (**auto.master** by default) which references the new indirect map. Because map nesting is only allowed in the **auto.master** map, a submount map referenced in other maps needs to follow a standard submount format (that is, *<key> <origin> <mapname>*) so that the referenced map is correctly loaded from LDAP. However, the **automountmap-add-indirect** sub-command did not follow this distinction and the *<origin>* and *<mapname>* attributes were not filled correctly. Therefore, submount maps referenced in a non-**auto.master** map were not recognized as automount maps by the **autofs** client software, and were not mounted. Submount maps referenced in a map that is not an **auto.master** map now follow a standard submount format, with the correct *<key>*, *<origin>* (**-fstype=autofs**), and *<mapname>* (**ldap:$MAP_NAME**). **autofs** client software is now able to correctly process submount maps both in auto.master and in other maps, and mount them.

BZ#785756

Prior to this update, the Identity Management user plug-in used a hard-coded default value for user's home directory instead of using the value that was configured. When an administrator changed the default user home directory in the Identity Management config plug-in from the default value to a custom value, this value was not honored when a user was added. This update fixes this bug, and when a new user is created without a custom home directory specified via a special option, the default configured home directory is used.

BZ#797274

The Identity Management certificate template did not include a **subjectKeyIdentifier** field even though it is marked with the *SHOULD* keyword in the RFC 3280 document. Because of this, certain applications processing these certificates could report errors. With this update, the certificate template for both current and new IPA server installations now contain the **subjectKeyIdentifier** field.

BZ#797562

Identity Management host and DNS plug-ins did not properly process hostnames or DNS zone names with a trailing dot. Consequently, the created host record FQDN attribute contained two values instead of one normalized value. This may have caused issues in further host record processing. With this update, all hostnames are normalized using a format without a trailing dot. The Identity Management DNS plug-in now accepts DNS zone names in both formats – with and without a trailing dot.

BZ#797565

Previously, CSVs were split in both CLI and server part of Identity Management processing. As a result, values which contained escaped comma characters were incorrectly split for the second time. With this update, CSV processing is done only in the client interface. Identity Management RPC interfaces (both XML-RPC and JSON-RPC) no longer process CSVs. Comma escaping was also replaced with quoting.

BZ#797566

The Identity Management server uninstall process removed system users that were added as a part of an Identity Management installation. This included **dirsrv** or **pkiuser** users, which the Directory Server uses to run its instances. These users also own log files produced by the Directory Server. If an Identity Management server was installed again, and the newly added system users' UIDs changed, the Directory Server could fail to start because the Directory Server instance was not permitted to write to the log files owned by the old system users with different UIDs. With this update, system users generated by an Identity Management server installation are no longer removed during the uninstall process.

BZ#747693

Identity Management plug-ins for LDAP ACI management (permission, selfservice, and delegation plug-ins) did not process their options in a robust way and had a relaxed validation of passed values. ACI management plug-ins could return *Internal Errors* when empty options or the **--raw** option were passed. An Internal Error was also returned when an invalid attribute was passed to the ACI attribute list option. Option processing is now more robust and more strict in validation. Proper errors are now returned when invalid or empty option values are passed.

BZ#746805

Objects which have an enabled/disabled state (that is, user accounts, sudo rules, HBAC rules, SELinux policies) were not distinguished in related search pages in the Web UI. Lines containing disabled objects are now grayed out in the search pages, and enabled columns have a different icon for each state.

BZ#802912

An Identity Management certificate did not read a custom user certificate subject base when validating a new certificate issuer. When an Identity Management server is installed with a custom subject base, and does not use the default subject base, issuing new certificates in the Identity Management Certificate Authority may return invalid issuer errors. With this update, a custom user certificate subject base is always read before the certificate issuer is validated, and the aforementioned errors are no longer returned when certificates are issued.

BZ#803050

Clicking **Cancel** in an error dialog in the Web UI when an unexpected error, such as an internal server error, was received made the Web UI unusable because the error message replaced the page content. With this update, error messages have their own containers, which fixes the aforementioned issue.

BZ#803836

Identity Management did not configure its Directory Server instance to always keep its RootDSE available anonymously and decrypted. As a consequence, when a user changed the **nsslapd-minssf** attribute in the Directory Server instance configuration to increase security demands on the connection to the instance, some applications (for example, **SSSD**) may have stopped working as they could no longer read RootDSE anonymously. To fix this issue, Identity Management now sets the **nsslapd-minssf-exclude-rootdse** option in the Directory Server instance configuration. Users and applications can access RootDSE in an Identity Management Directory Server instance anonymously even when the instance is configured with increased security demands on incoming connections.

BZ#807366

Previously, the Netgroup page in the Web UI did not have input fields for specifying **all** options. With this update, the entire Netgroup page has been redesigned to add this functionality.

BZ#688765

Identity Management DNS plug-in did not validate the contents of DNS records. Some DNS record types (for example, MX, LOC, or SRV) have a complex data structure which needs to be stored, otherwise the record is not resolvable. Relaxed DNS plug-in validation let users create invalid records which then could not be resolved even though they were stored in LDAP. With this update, every DNS record type (except the experimental A6 DNS record type) is now validated with respect to a relevant RFC document. The validation covers most common user errors and also provides the user with guidance on why the entered record is invalid. Users are also able to create more complex DNS records without detailed knowledge of their structure as the improved DNS plug-in interface provides guidance when creating DNS records. Also, the DNS plug-in does not let users enter invalid records any more.

Enhancements



NOTE

For a list of major features that were added by this update, refer to Red Hat Enterprise Linux 6.3 Release Notes.

BZ#759501

When the number of failed login attempts exceeds the maximum that is configured, the account is locked. However, an investigation of the lock-out status of a particular user was difficult as the

number of failed login attempts was not replicated. Identity Management now includes a new **ipa user-status** command that provides the number of failed login attempts on all configured replicas along with the time of the last successful or failed login attempt.

BZ#766181

When a new user is added, a User Private Group (UPG) is created and assigned as that user's primary group by default. However, there may be use cases when an administrator wants to use a common group assigned as a primary group for all users. The Directory Server plug-in that handles the creation of UPGs can now be disabled with a new utility — **ipa-managed-entries**. This utility lets administrators disable automatic creation of UPGs, and allows all new users to share a common group as their primary group.

BZ#767725

When an Identity Management server is configured with DNS support, DNS zone dynamic update policy allows Identity Management clients to update a relevant DNS forward record if the client IP address changes. However, for security reasons, clients cannot be allowed to update their reverse records because they would be able to change any record in the reverse zone. With this update, an Identity Management DNS zone can be configured to allow automatic updates of client reverse records when the forward record is updated with the new IP address. As a result, both forward and reverse records for a client machine can be updated when the client IP address changes.

BZ#772044

The Identity Management **host** plug-in did not allow storing of machine MAC addresses. Administrators could not assign MAC addresses to host entries in Identity Management. With this update, a new attribute for MAC addresses was added to the Identity Management host plug-in. Administrators can now assign a MAC address to a host entry. The value can then be read from the Identity Management LDAP server with, for example, the following command:

```
~]$ getent ethers <hostname>
```

BZ#772301

When a forward DNS record was created, no corresponding reverse record was created even when both the forward and the reverse zone were managed by Identity Management. Users always had to create both the forward and the reverse records manually. With this update, both CLI and Web UI now have the option to automatically create a reverse record when an IPv4 or IPv6 forward record is created.

BZ#807361

Prior to this update, all DNS records in an Identity Management Directory Server instance were publicly accessible. With a publicly accessible DNS tree in the Directory Server instance, anyone with access to the server could acquire all DNS data. This operation is normally restricted with access control rules. It is a common security practice to keep this information restricted to only a selected group of users. Therefore, with this update, the entire LDAP tree with DNS records is now accessible only to the LDAP driver which feeds the data to the name server, admin users, or users with a new permission called **Read DNS Entries**. As a result, only permitted users can now access all DNS records in Identity Management Directory Server instances.

BZ#753483

The Identity Management server did not allow the creation of DNS zones with conditional forwarding, which lets the name server forward all zone requests to a custom forwarder. With this update, the Identity Management DNS plug-in allows users to create a DNS zone and set a conditional forwarder and a forwarding policy for that zone.

BZ#803822

Support for SSH public key management was added to Identity Management server; OpenSSH on Identity Management clients is automatically configured to use the public keys stored on the Identity Management server. This feature is a Technology Preview.

BZ#745968

The DNS page in the Web UI did not allow navigation from A or AAAA records to the related PTR records. This update adds a link which points to a related PTR record if it exists.

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

5.111. IPMITOOL

5.111.1. RHBA-2013:0264 – ipmitool bug fix update

An updated ipmitool package that fixes one bug is now available for Red Hat Enterprise Linux 6 Extended Update Support.

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 Fix

BZ#907926

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.

All users of ipmitool are advised to upgrade to this updated package, which fixes this bug.

5.111.2. RHBA-2012:0999 – ipmitool bug fix update

An updated ipmitool package that fixes one bug is now available for Red Hat Enterprise Linux 6.

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 Fix

BZ#828678

In the previous ipmitool package update, 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 wrong default value of the retransmission timeout and outgoing request timed out prematurely. In addition, in some corner cases, ipmitool could have terminated unexpectedly with a segmentation fault when the timeout occurred. This update fixes the default timeout value and ipmitool without the "-N" option retransmits outgoing IPMI requests like in previous versions.

All users of ipmitool are advised to upgrade to this updated package, which fixes this bug.

5.111.3. RHBA-2012:0875 – ipmitool bug fix and enhancement update

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

The ipmitool packages contain 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#715615

Previously, the exit code of the "ipmitool -o list" command was set incorrectly so that the command always returned 1. This update modifies ipmitool to return the exit code 0 as expected.

BZ#725993

The "ipmitool sol payload" and "ipmitool sel" commands previously accepted incorrect argument values, which caused the ipmitool utility to terminate unexpectedly with a segmentation fault. With this update, argument values of these commands are now validated, and ipmitool no longer crashes but generates an error message when used with incorrect arguments.

Enhancements

BZ#748073

Previously, ipmitool could not be used to set retransmission intervals of IPMI messages over the LAN or lanplus interface. This update introduces new options, "-R" and "-N", which can be used to specify number of retransmissions and delay between them (in seconds) when transferring IPMI messages using the LAN or lanplus interfaces.

BZ#739358

The "ipmitool delloem" command has been updated to the latest upstream version, which includes the new "vFlash" command allowing to show information about extended SD cards. This patch also updates documentation of the "ipmitool delloem" commands, improves error descriptions and adds support for new hardware.

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

5.112. IPROUTE

5.112.1. RHBA-2012:0835 – iproute bug fix and enhancement update

An updated iproute package that fixes two bugs and adds three enhancements is now available for Red Hat Enterprise Linux 6.

The iproute package contains networking utilities (ip and rtmon, for example), which are designed to use the advanced networking capabilities of the Linux kernel.

Bug Fixes

BZ#730627

The ip6tunnel mode command passed a zeroed parameter structure to the kernel, which attempted to change all tunnel parameters to zero and failed. Consequently, users could not change ip6tunnel parameters. With this update, the ip6tunnel code has been changed so that it updates only the changed parameters. As a result, it is now possible for users to adjust ip6tunnel parameters as expected.

BZ#736106

The lnstat utility used an incorrect file descriptor for its dump output. Consequently, the lnstat utility printed its dump output to stderr rather than to stdout. The code has been fixed and lnstat now prints its dump output to stdout.

Enhancements

BZ#748767

The tc utility (a traffic control tool) has been enhanced to allow users to work with the Multi-queue priority (MQPRIO) Queueing Discipline (qdiscs) scheduler. With MQPRIO qdiscs, QOS can be offloaded from NICs that support external QOS schedulers. As a result, it is now possible for users to monitor traffic classes, gather statistics, set socket-buffer (SKB) priority and socket-priority-to-traffic-class mapping.

BZ#788120

The tc utility has been updated to work with Quick Fair Queueing (QFQ) kernel features. Users can now take advantage of the new QFQ-traffic scheduler from user space.

BZ#812779

This update adds support for multiple multicast routing tables.

Users are advised to upgrade to this updated iproute package, which fixes these bugs and adds these enhancements.

5.113. IPRUTILS

5.113.1. RHBA-2012:1183 – iprutils bug fix update

An updated iprutils package that fixes a bug now available for Red Hat Enterprise Linux 6.

The iprutils package provides utilities to manage and configure SCSI devices that are supported by the IBM Power RAID SCSI storage device driver.

Bug Fix

BZ#849556

Previously, a buffer overflow bug caused the iprconfig utility to terminate unexpectedly with a segmentation fault when displaying detailed information of a disk device. A patch has been provided to address this issue and iprconfig no longer crashes in the described scenario.

All users of iprutils are advised to upgrade to this updated package, which fixes this bug.

5.113.2. RHBA-2012:0792 – iprutils bug fix and enhancement update

Updated iprutils packages that fix multiple bugs add various enhancements is now available for Red Hat Enterprise Linux 6.

The iprutils package provides utilities to manage and configure SCSI devices that are supported by the IBM Power RAID SCSI storage device driver.

The iprutils package has been upgraded to upstream version 2.3.9, which fixes multiple bugs and adds multiple enhancements. These packages also add support for the CRoC-based 6 GB Serial Attached SCSI (SAS) vRAID adapters on IBM POWER7. (BZ#738890, BZ#817087)

All users of iprutils are advised to upgrade to these updated packages, which add fix these enhancements and add these enhancements.

5.114. IPTRAF

5.114.1. RHBA-2012:0762 – iptraf bug fix update

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

IPTraf is a console-based network monitoring utility. IPTraf gathers data such as TCP connection packet and byte counts, interface statistics and activity indicators, TCP/UDP traffic breakdowns, and LAN station packet and byte counts.

Bug Fix

BZ#682350

Prior to this update, interface names were checked by IPTraf against a whitelist of names to determine whether an interface was supported. Network devices can have arbitrary names and due to the changes for "Consistent Network Device Naming", the interface names will change to location-based names. Consequently, IPTraf could reject certain interface names. This update removes the interface name check and as a result IPTraf always accepts device names.

All users of iptraf are advised to upgrade to these updated packages, which fix this bug.

5.115. IPVSADM

5.115.1. RHBA-2012:0865 – ipvsadm bug fix update

Updated ipvsadm packages that fix one bug is now available for Red Hat Enterprise Linux 6.

The ipvsadm package provides the ipsvadm tool to administer the IP Virtual Server services offered by the Linux kernel.

Bug Fix

BZ#788529

Prior to this update, the ipvsadm utility did not correctly handle out-of-order messages from the kernel concerning the sync daemon. As a consequence, the "ipvsadm --list --daemon" command did not always output the status of the sync daemon. With this update, the ordering of messages from the kernel no longer influences the output, and the command always returns the sync daemon status.

All users of ipvsadm are advised to upgrade to these updated packages, which fix this bug.

5.116. IRQBALANCE

5.116.1. RHBA-2012:1157 – irqbalance bug fix update

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

The irqbalance package provides a daemon that evenly distributes interrupt request (IRQ) load across multiple CPUs for enhanced performance.

Bug Fix

BZ#845374

The irqbalance daemon assigns each interrupt source in the system to a "class", which represents the type of the device (for example Networking, Storage or Media). Previously, irqbalance had some problems while classifying certain NIC devices that resulted into performance impact on affected systems. With this update, the NIC classification mechanism has been updated to work with all types of NICs.

All users of irqbalance are advised to upgrade to these updated packages, which fix this bug.

5.116.2. RHBA-2012:0807 – irqbalance bug fix update

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

The irqbalance packages provide a daemon that evenly distributes interrupt requests (IRQ) load across multiple CPUs for enhanced performance.

Bug Fix

BZ#682211

The irqbalance daemon assigns each interrupt source in the system to a "class", which represents the type of the device (for example Networking, Storage or Media). Previously, irqbalance used the IRQ handler names from the /proc/interrupts file to decide the source class, which caused irqbalance to not recognize network interrupts correctly. As a consequence, systems using biosdevname NIC naming did not have their hardware interrupts distributed and pinned as expected. With this update, the device classification mechanism has been improved, and so ensures a better interrupts distribution.

All users of irqbalance are advised to upgrade to these updated packages, which fix this bug.

5.117. IRSSI

5.117.1. RHBA-2012:1171 – irssi bug fix update

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

Irssi is a modular IRC client with Perl scripting. Only the text-mode front end is currently supported.

Bug Fixes

BZ#639258

Prior to this update, when the user attempted to use the "/unload" command to unload a static module, Irssi incorrectly marked this module as unavailable, rendering the user unable to load this module again without restarting the client. This update adapts the underlying source code to ensure that only dynamic modules can be unloaded.

BZ#845047

The previous version of the irssi(1) manual page documented "--usage" as a valid command line option. This was incorrect, because Irssi no longer supports this option and an attempt to use it causes it to fail with an error. With this update, the manual page has been corrected and no longer documents unsupported command line options.

All users of irssi are advised to upgrade to these updated packages, which fix these bugs.

5.118. ISCSI-INITIATOR-UTILS

5.118.1. RHBA-2012:0957 – 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 package provides the server daemon for the iSCSI protocol, as well as utilities used to manage the daemon. iSCSI is a protocol for distributed disk access using SCSI commands sent over Internet Protocol networks.

The iscsiuio tool has been upgraded to upstream version 0.7.2.1, which provides a number of bug fixes and one enhancement over the previous version. (BZ#740054)

Bug Fixes

BZ#738192

The iscsistart utility used hard-coded values as its settings. Consequently, it could take several minutes before change failure detection and path failover when using dm-multipath took place. With this update, the iscsistart utility has been modified to process settings provided on the command line.

BZ#739049

The iSCSI README file incorrectly listed the --info option as the option to display iscsiadm iSCSI information. The README has been corrected and it now states correctly that you need to use the "-P 1" argument to obtain such information.

BZ#739843

The iSCSI discovery process via a TOE (TCP Offload Engine) interface failed if the "iscsiadm -m iface" command had not been executed. This happened because the "iscsiadm -m" discovery command did not check interface settings. With this update, the iscsiadm tool creates the default ifaces settings when first used and the problem no longer occurs.

BZ#796574

If the port number was passed with a non-fully-qualified hostname to the iscsiadm tool, the tool created records with the port being part of the hostname. Consequently, the login or discovery operation failed because iscsiadm was not able to find the record. With this update, the iscsiadm

portal parser has been modified to separate the port from the hostname. As a result, the port is parsed and processed correctly.

Enhancement

BZ#790609

The iscsidm tool has been updated to support the ping command using QLogic's iSCSI offload cards and to manage the CHAP (Challenge-Handshake Authentication Protocol) entries on the host.

All users of iscsi-initiator-utils are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.119. JAKARTA-COMMONS-HTTPCLIENT

5.119.1. RHSA-2013:0270 – Moderate: jakarta-commons-httpclient security update

Updated jakarta-commons-httpclient packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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

The Jakarta Commons HttpClient component can be used to build HTTP-aware client applications (such as web browsers and web service clients).

Security Fix

CVE-2012-5783

The Jakarta Commons HttpClient component did not verify that the server hostname matched the domain name in the subject's Common Name (CN) or subjectAltName field in X.509 certificates. This could allow a man-in-the-middle attacker to spoof an SSL server if they had a certificate that was valid for any domain name.

All users of jakarta-commons-httpclient are advised to upgrade to these updated packages, which correct this issue. Applications using the Jakarta Commons HttpClient component must be restarted for this update to take effect.

5.120. JAVA-1.5.0-IBM

5.120.1. RHSA-2012:1465 – Critical: java-1.5.0-ibm security update

Updated java-1.5.0-ibm packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

IBM J2SE version 5.0 includes the IBM Java Runtime Environment and the IBM Java Software Development Kit.

Security Fix

CVE-2012-1531, CVE-2012-3143, CVE-2012-3216, CVE-2012-4820, CVE-2012-4822, CVE-2012-5069, CVE-2012-5071, CVE-2012-5073, CVE-2012-5075, CVE-2012-5079, CVE-2012-5081, CVE-2012-5083, CVE-2012-5084, CVE-2012-5089

This update fixes several vulnerabilities in the IBM Java Runtime Environment and the IBM Java Software Development Kit. Detailed vulnerability descriptions are linked from the IBM Security alerts page.

All users of java-1.5.0-ibm are advised to upgrade to these updated packages, containing the IBM J2SE 5.0 SR15 release. All running instances of IBM Java must be restarted for this update to take effect.

5.120.2. RHSA-2012:1245 — Critical: java-1.5.0-ibm security update

Updated java-1.5.0-ibm packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

IBM J2SE version 5.0 includes the IBM Java Runtime Environment and the IBM Java Software Development Kit.

Security Fix

CVE-2012-1713, CVE-2012-1716, CVE-2012-1717, CVE-2012-1718, CVE-2012-1719, CVE-2012-1725

This update fixes several vulnerabilities in the IBM Java Runtime Environment and the IBM Java Software Development Kit. Detailed vulnerability descriptions are linked from the IBM Security alerts page.

All users of java-1.5.0-ibm are advised to upgrade to these updated packages, containing the IBM J2SE 5.0 SR14 release. All running instances of IBM Java must be restarted for this update to take effect.

5.121. JAVA-1.6.0-IBM

5.121.1. RHSA-2012:1466 — Critical: java-1.6.0-ibm security update

Updated java-1.6.0-ibm packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

IBM Java SE version 6 includes the IBM Java Runtime Environment and the IBM Java Software Development Kit.

Security Fix

CVE-2012-0547, CVE-2012-1531, CVE-2012-1532, CVE-2012-1533, CVE-2012-1682, CVE-2012-3143, CVE-2012-3159, CVE-2012-3216, CVE-2012-4820, CVE-2012-4822, CVE-2012-4823, CVE-

CVE-2012-5068, CVE-2012-5069, CVE-2012-5071, CVE-2012-5072, CVE-2012-5073, CVE-2012-5075, CVE-2012-5079, CVE-2012-5081, CVE-2012-5083, CVE-2012-5084, CVE-2012-5089

This update fixes several vulnerabilities in the IBM Java Runtime Environment and the IBM Java Software Development Kit. Detailed vulnerability descriptions are linked from the IBM Security alerts page.

All users of java-1.6.0-ibm are advised to upgrade to these updated packages, containing the IBM Java SE 6 SR12 release. All running instances of IBM Java must be restarted for the update to take effect.

5.121.2. RHSA-2012:1238 — Critical: java-1.6.0-ibm security update

Updated java-1.6.0-ibm packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

IBM Java SE version 6 includes the IBM Java Runtime Environment and the IBM Java Software Development Kit.

Security Fix

CVE-2012-0551, CVE-2012-1713, CVE-2012-1716, CVE-2012-1717, CVE-2012-1718, CVE-2012-1719, CVE-2012-1721, CVE-2012-1722, CVE-2012-1725

This update fixes several vulnerabilities in the IBM Java Runtime Environment and the IBM Java Software Development Kit. Detailed vulnerability descriptions are linked from the IBM Security alerts page.

All users of java-1.6.0-ibm are advised to upgrade to these updated packages, containing the IBM Java SE 6 SR11 release. All running instances of IBM Java must be restarted for the update to take effect.

5.122. JAVA-1.6.0-OPENJDK

5.122.1. RHSA-2012:1384 — Critical: java-1.6.0-openjdk security update

Updated java-1.6.0-openjdk packages that fix several security issues are now available for Red Hat Enterprise Linux 6.

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

These packages provide the OpenJDK 6 Java Runtime Environment and the OpenJDK 6 Software Development Kit.

Security Fixes

CVE-2012-5086, CVE-2012-5084, CVE-2012-5089

Multiple improper permission check issues were discovered in the Beans, Swing, and JMX components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

CVE-2012-5068, CVE-2012-5071, CVE-2012-5069, CVE-2012-5073, CVE-2012-5072

Multiple improper permission check issues were discovered in the Scripting, JMX, Concurrency, Libraries, and Security components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass certain Java sandbox restrictions.

CVE-2012-5079

It was discovered that java.util.ServiceLoader could create an instance of an incompatible class while performing provider lookup. An untrusted Java application or applet could use this flaw to bypass certain Java sandbox restrictions.

CVE-2012-5081

It was discovered that the Java Secure Socket Extension (JSSE) SSL/TLS implementation did not properly handle handshake records containing an overly large data length value. An unauthenticated, remote attacker could possibly use this flaw to cause an SSL/TLS server to terminate with an exception.

CVE-2012-5075

It was discovered that the JMX component in OpenJDK could perform certain actions in an insecure manner. An untrusted Java application or applet could possibly use this flaw to disclose sensitive information.

CVE-2012-4416

A bug in the Java HotSpot Virtual Machine optimization code could cause it to not perform array initialization in certain cases. An untrusted Java application or applet could use this flaw to disclose portions of the virtual machine's memory.

CVE-2012-5077

It was discovered that the SecureRandom class did not properly protect against the creation of multiple seeders. An untrusted Java application or applet could possibly use this flaw to disclose sensitive information.

CVE-2012-3216

It was discovered that the java.io.FilePermission class exposed the hash code of the canonicalized path name. An untrusted Java application or applet could possibly use this flaw to determine certain system paths, such as the current working directory.

CVE-2012-5085

This update disables Gopher protocol support in the java.net package by default. Gopher support can be enabled by setting the newly introduced property, "jdk.net.registerGopherProtocol", to true.



NOTE

If the web browser plug-in provided by the icedtea-web package was installed, the issues exposed via Java applets could have been exploited without user interaction if a user visited a malicious website.

This erratum also upgrades the OpenJDK package to IcedTea6 1.11.5. Refer to the NEWS file for further information:

http://icedtea.classpath.org/hg/release/icedtea6-1.11/file/icedtea6-1.11.5/NEWS

All users of java-1.6.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.122.2. RHSA-2013:0245 – Critical: java-1.6.0-openjdk security update

Updated java-1.6.0-openjdk packages that fix several security issues are now available for Red Hat Enterprise Linux 6.

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

These packages provide the OpenJDK 6 Java Runtime Environment and the OpenJDK 6 Software Development Kit.

Security Fixes

CVE-2013-0442, CVE-2013-0445, CVE-2013-0441, CVE-2013-1475, CVE-2013-1476, CVE-2013-0429, CVE-2013-0450, CVE-2013-0425, CVE-2013-0426, CVE-2013-0428

Multiple improper permission check issues were discovered in the AWT, CORBA, JMX, and Libraries components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

CVE-2013-1478, CVE-2013-1480

Multiple flaws were found in the way image parsers in the 2D and AWT components handled image raster parameters. A specially-crafted image could cause Java Virtual Machine memory corruption and, possibly, lead to arbitrary code execution with the virtual machine privileges.

CVE-2013-0432

A flaw was found in the AWT component's clipboard handling code. An untrusted Java application or applet could use this flaw to access clipboard data, bypassing Java sandbox restrictions.

CVE-2013-0435

The default Java security properties configuration did not restrict access to certain com.sun.xml.internal packages. An untrusted Java application or applet could use this flaw to access information, bypassing certain Java sandbox restrictions. This update lists the whole package as restricted.

CVE-2013-0427, CVE-2013-0433, CVE-2013-0434

Multiple improper permission check issues were discovered in the Libraries, Networking, and JAXP components. An untrusted Java application or applet could use these flaws to bypass certain Java sandbox restrictions.

CVE-2013-0424

It was discovered that the RMI component's CGIHandler class used user inputs in error messages without any sanitization. An attacker could use this flaw to perform a cross-site scripting (XSS) attack.

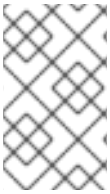
CVE-2013-0440

It was discovered that the SSL/TLS implementation in the JSSE component did not properly

enforce handshake message ordering, allowing an unlimited number of handshake restarts. A remote attacker could use this flaw to make an SSL/TLS server using JSSE consume an excessive amount of CPU by continuously restarting the handshake.

CVE-2013-0443

It was discovered that the JSSE component did not properly validate Diffie-Hellman public keys. An SSL/TLS client could possibly use this flaw to perform a small subgroup attack.



NOTE

If the web browser plug-in provided by the icedtea-web package was installed, the issues exposed via Java applets could have been exploited without user interaction if a user visited a malicious website.

This erratum also upgrades the OpenJDK package to IcedTea6 1.11.6. Refer to the NEWS file for further information:

http://icedtea.classpath.org/hg/release/icedtea6-1.11/file/icedtea6-1.11.6/NEWS

All users of java-1.6.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.122.3. RHSA-2012:1221 – Critical: java-1.6.0-openjdk security update

Updated java-1.6.0-openjdk packages that fix two security issues are now available for Red Hat Enterprise Linux 6.

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

These packages provide the OpenJDK 6 Java Runtime Environment and the OpenJDK 6 Software Development Kit.

Security Fixes

CVE-2012-1682

It was discovered that the Beans component in OpenJDK did not perform permission checks properly. An untrusted Java application or applet could use this flaw to use classes from restricted packages, allowing it to bypass Java sandbox restrictions.

CVE-2012-0547

A hardening fix was applied to the AWT component in OpenJDK, removing functionality from the restricted SunToolkit class that was used in combination with other flaws to bypass Java sandbox restrictions.

Note: If the web browser plug-in provided by the icedtea-web package was installed, the issues exposed via Java applets could have been exploited without user interaction if a user visited a malicious website.

This erratum also upgrades the OpenJDK package to IcedTea6 1.11.4. Refer to the NEWS file for further information.

All users of java-1.6.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.122.4. RHSA-2013:0273 — Critical: java-1.6.0-openjdk security update

Updated java-1.6.0-openjdk packages that fix two security issues are now available for Red Hat Enterprise Linux 6.

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

These packages provide the OpenJDK 6 Java Runtime Environment and the OpenJDK 6 Software Development Kit.

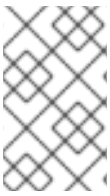
Security Fixes

CVE-2013-1486

An improper permission check issue was discovered in the JMX component in OpenJDK. An untrusted Java application or applet could use this flaw to bypass Java sandbox restrictions.

CVE-2013-0169

It was discovered that OpenJDK leaked timing information when decrypting TLS/SSL protocol encrypted records when CBC-mode cipher suites were used. A remote attacker could possibly use this flaw to retrieve plain text from the encrypted packets by using a TLS/SSL server as a padding oracle.



NOTE

If the web browser plug-in provided by the icedtea-web package was installed, CVE-2013-1486 could have been exploited without user interaction if a user visited a malicious website.

This erratum also upgrades the OpenJDK package to IcedTea6 1.11.8. Refer to the NEWS file for further information:

http://icedtea.classpath.org/hg/release/icedtea6-1.11/file/icedtea6-1.11.8/NEWS

All users of java-1.6.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.122.5. RHBA-2012:0836 — java-1.6.0-openjdk bug fix and enhancement update

Updated java-1.6.0-openjdk packages that fix multiple 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 Software Development Kit.

The java-1.6.0-openjdk packages have been upgraded to upstream version 1.11.1, which provides a number of bug fixes and enhancements over the previous version. (BZ#771971)

Bug Fixes

BZ#751203

Previously, after updating OpenJDK to java-1.6.0-openjdk-1.6.0.0-1.40.1.9.10.el6_1, the Java Remote Object Registry (rmiregistry) started only if run with the java.rmi.server.codebase argument, otherwise the registry start failed. This update fixes the regression and the registry can be started without the argument as expected.

BZ#767537

Channel binding for the Kerberos protocol was implemented incorrectly and OpenJDK did not process Kerberos GSS (General Security Services) contexts which did not have incoming channel binding. This resulted in interopability problems with Internet Explorer on Windows Server 2008. With this update, OpenJDK handles unset channel binding correctly and processes Kerberos GSS contexts as expected.

BZ#804632

The SystemTap script translator (stap) run with jstack() systemtap support could terminate with an error similar to the following:

```
ERROR: kernel read fault at 0x0000000000000018 (addr) near identifier '@cast' at
/usr/share/systemtap/tapset/x86_64/jstack.stp:362:29
```

This update improves the jstack code including, for example, the constant definition and error handling, and the stap script with jstack now works more reliably.

BZ#805936, BZ#807324

This update fixes multiple problems that occurred when using signed jar files.

Enhancement

BZ#751410

Support for huge pages was added.

All users of java-1.6.0-openjdk are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.123. JAVA-1.6.0-SUN

5.123.1. RHSA-2013:0236 – Critical: java-1.6.0-sun security update

Updated java-1.6.0-sun packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

Oracle Java SE version 6 includes the Oracle Java Runtime Environment and the Oracle Java Software Development Kit.

Security Fix

CVE-2012-1541, CVE-2012-3213, CVE-2012-3342, CVE-2013-0351, CVE-2013-0409, CVE-2013-0419, CVE-2013-0423, CVE-2013-0424, CVE-2013-0425, CVE-2013-0426, CVE-2013-0427, CVE-2013-0428, CVE-2013-0429, CVE-2013-0430, CVE-2013-0432, CVE-2013-0433, CVE-2013-0434, CVE-2013-0435, CVE-2013-0438, CVE-2013-0440, CVE-2013-0441, CVE-2013-0442, CVE-2013-0443, CVE-2013-0445, CVE-2013-0446, CVE-2013-0450, CVE-2013-1473, CVE-2013-1475, CVE-2013-1476, CVE-2013-1478, CVE-2013-1480, CVE-2013-1481

This update fixes several vulnerabilities in the Oracle Java Runtime Environment and the Oracle Java Software Development Kit. Further information about these flaws can be found on the Oracle Java SE Critical Patch Update Advisory page.

All users of java-1.6.0-sun are advised to upgrade to these updated packages, which provide Oracle Java 6 Update 39. All running instances of Oracle Java must be restarted for the update to take effect.

5.123.2. RHSA-2012:1392 – Critical: java-1.6.0-sun security update

Updated java-1.6.0-sun packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

Oracle Java SE version 6 includes the Oracle Java Runtime Environment and the Oracle Java Software Development Kit.

Security Fix

CVE-2012-0547, CVE-2012-1531, CVE-2012-1532, CVE-2012-1533, CVE-2012-3143, CVE-2012-3159, CVE-2012-3216, CVE-2012-4416, CVE-2012-5068, CVE-2012-5069, CVE-2012-5071, CVE-2012-5072, CVE-2012-5073, CVE-2012-5075, CVE-2012-5077, CVE-2012-5079, CVE-2012-5081, CVE-2012-5083, CVE-2012-5084, CVE-2012-5085, CVE-2012-5086, CVE-2012-5089

This update fixes several vulnerabilities in the Oracle Java Runtime Environment and the Oracle Java Software Development Kit. Further information about these flaws can be found on the Oracle Java SE Critical Patch Update Advisory and Oracle Security Alert pages .

All users of java-1.6.0-sun are advised to upgrade to these updated packages, which provide Oracle Java 6 Update 37. All running instances of Oracle Java must be restarted for the update to take effect.

5.124. JAVA-1.7.0-IBM

5.124.1. RHSA-2012:1467 – Critical: java-1.7.0-ibm security update

Updated java-1.7.0-ibm packages that fix several security issues are now available for Red Hat Enterprise Linux 6 Supplementary.

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

IBM Java SE version 7 includes the IBM Java Runtime Environment and the IBM Java Software Development Kit.

Security Fix

CVE-2012-1531, CVE-2012-1532, CVE-2012-1533, CVE-2012-1718, CVE-2012-3143, CVE-2012-3159, CVE-2012-3216, CVE-2012-4820, CVE-2012-4821, CVE-2012-4822, CVE-2012-4823, CVE-2012-5067, CVE-2012-5069, CVE-2012-5070, CVE-2012-5071, CVE-2012-5072, CVE-2012-5073, CVE-2012-5074, CVE-2012-5075, CVE-2012-5076, CVE-2012-5077, CVE-2012-5079, CVE-2012-5081, CVE-2012-5083, CVE-2012-5084, CVE-2012-5086, CVE-2012-5087, CVE-2012-5088, CVE-2012-5089

This update fixes several vulnerabilities in the IBM Java Runtime Environment and the IBM Java Software Development Kit. Detailed vulnerability descriptions are linked from the IBM Security alerts page.

All users of java-1.7.0-ibm are advised to upgrade to these updated packages, containing the IBM Java SE 7 SR3 release. All running instances of IBM Java must be restarted for the update to take effect.

5.124.2. RHSA-2012:1289 — Critical: java-1.7.0-ibm security update

Updated java-1.7.0-ibm packages that fix several security issues are now available for Red Hat Enterprise Linux 6 Supplementary.

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

IBM Java SE version 7 includes the IBM Java Runtime Environment and the IBM Java Software Development Kit.

Security Fix

CVE-2012-0547, CVE-2012-0551, CVE-2012-1682, CVE-2012-1713, CVE-2012-1716, CVE-2012-1717, CVE-2012-1719, CVE-2012-1721, CVE-2012-1722, CVE-2012-1725, CVE-2012-1726, CVE-2012-3136, CVE-2012-4681

This update fixes several vulnerabilities in the IBM Java Runtime Environment and the IBM Java Software Development Kit. Detailed vulnerability descriptions are linked from the IBM Security alerts page.

All users of java-1.7.0-ibm are advised to upgrade to these updated packages, containing the IBM Java SE 7 SR2 release. All running instances of IBM Java must be restarted for the update to take effect.

5.125. JAVA-1.7.0-OPENJDK

5.125.1. RHSA-2013:0247 — Important: java-1.7.0-openjdk security update

Updated java-1.7.0-openjdk packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

These packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Software Development Kit.

Security Fixes

CVE-2013-0442, CVE-2013-0445, CVE-2013-0441, CVE-2013-1475, CVE-2013-1476, CVE-2013-0429, CVE-2013-0450, CVE-2013-0425, CVE-2013-0426, CVE-2013-0428, CVE-2013-0444

Multiple improper permission check issues were discovered in the AWT, CORBA, JMX, Libraries, and Beans components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

CVE-2013-1478, CVE-2013-1480

Multiple flaws were found in the way image parsers in the 2D and AWT components handled image raster parameters. A specially-crafted image could cause Java Virtual Machine memory corruption and, possibly, lead to arbitrary code execution with the virtual machine privileges.

CVE-2013-0432

A flaw was found in the AWT component's clipboard handling code. An untrusted Java application or applet could use this flaw to access clipboard data, bypassing Java sandbox restrictions.

CVE-2013-0435

The default Java security properties configuration did not restrict access to certain com.sun.xml.internal packages. An untrusted Java application or applet could use this flaw to access information, bypassing certain Java sandbox restrictions. This update lists the whole package as restricted.

CVE-2013-0431, CVE-2013-0427, CVE-2013-0433, CVE-2013-0434

Multiple improper permission check issues were discovered in the JMX, Libraries, Networking, and JAXP components. An untrusted Java application or applet could use these flaws to bypass certain Java sandbox restrictions.

CVE-2013-0424

It was discovered that the RMI component's CGIHandler class used user inputs in error messages without any sanitization. An attacker could use this flaw to perform a cross-site scripting (XSS) attack.

CVE-2013-0440

It was discovered that the SSL/TLS implementation in the JSSE component did not properly enforce handshake message ordering, allowing an unlimited number of handshake restarts. A remote attacker could use this flaw to make an SSL/TLS server using JSSE consume an excessive amount of CPU by continuously restarting the handshake.

CVE-2013-0443

It was discovered that the JSSE component did not properly validate Diffie-Hellman public keys. An SSL/TLS client could possibly use this flaw to perform a small subgroup attack.

This erratum also upgrades the OpenJDK package to IcedTea7 2.3.5.

All users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.125.2. RHSA-2013:0275 – Important: java-1.7.0-openjdk security update

Updated java-1.7.0-openjdk packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

These packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Software Development Kit.

Security Fixes

CVE-2013-1486, CVE-2013-1484

Multiple improper permission check issues were discovered in the JMX and Libraries components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

CVE-2013-1485

An improper permission check issue was discovered in the Libraries component in OpenJDK. An untrusted Java application or applet could use this flaw to bypass certain Java sandbox restrictions.

CVE-2013-0169

It was discovered that OpenJDK leaked timing information when decrypting TLS/SSL protocol encrypted records when CBC-mode cipher suites were used. A remote attacker could possibly use this flaw to retrieve plain text from the encrypted packets by using a TLS/SSL server as a padding oracle.

This erratum also upgrades the OpenJDK package to IcedTea7 2.3.7. Refer to the NEWS file for further information:

http://icedtea.classpath.org/hg/release/icedtea7-2.3/file/icedtea-2.3.7/NEWS

All users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.125.3. RHSA-2012:1386 – Important: java-1.7.0-openjdk security update

Updated java-1.7.0-openjdk packages that fix several security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

[Update 13 November 2012] The file list of this advisory was updated to move java-1.7.0-openjdk-devel from the optional repositories to the base repositories. Additionally, java-1.7.0-openjdk for the HPC Node variant was also moved (this package was already in the base repositories for other product variants). No changes have been made to the packages themselves.

These packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Software Development Kit.

Security Fixes

CVE-2012-5086, CVE-2012-5087, CVE-2012-5088, CVE-2012-5084, CVE-2012-5089

Multiple improper permission check issues were discovered in the Beans, Libraries, Swing, and JMX components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

CVE-2012-5076, CVE-2012-5074

The default Java security properties configuration did not restrict access to certain com.sun.org.glassfish packages. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions. This update lists those packages as restricted.

CVE-2012-5068, CVE-2012-5071, CVE-2012-5069, CVE-2012-5073, CVE-2012-5072

Multiple improper permission check issues were discovered in the Scripting, JMX, Concurrency, Libraries, and Security components in OpenJDK. An untrusted Java application or applet could use these flaws to bypass certain Java sandbox restrictions.

CVE-2012-5079

It was discovered that java.util.ServiceLoader could create an instance of an incompatible class while performing provider lookup. An untrusted Java application or applet could use this flaw to bypass certain Java sandbox restrictions.

CVE-2012-5081

It was discovered that the Java Secure Socket Extension (JSSE) SSL/TLS implementation did not properly handle handshake records containing an overly large data length value. An unauthenticated, remote attacker could possibly use this flaw to cause an SSL/TLS server to terminate with an exception.

CVE-2012-5070, CVE-2012-5075

It was discovered that the JMX component in OpenJDK could perform certain actions in an insecure manner. An untrusted Java application or applet could possibly use these flaws to disclose sensitive information.

CVE-2012-4416

A bug in the Java HotSpot Virtual Machine optimization code could cause it to not perform array initialization in certain cases. An untrusted Java application or applet could use this flaw to disclose portions of the virtual machine's memory.

CVE-2012-5077

It was discovered that the SecureRandom class did not properly protect against the creation of multiple seeders. An untrusted Java application or applet could possibly use this flaw to disclose sensitive information.

CVE-2012-3216

It was discovered that the java.io.FilePermission class exposed the hash code of the canonicalized path name. An untrusted Java application or applet could possibly use this flaw to determine certain system paths, such as the current working directory.

CVE-2012-5085

This update disables Gopher protocol support in the java.net package by default. Gopher support can be enabled by setting the newly introduced property, "jdk.net.registerGopherProtocol", to true.

This erratum also upgrades the OpenJDK package to IcedTea7 2.3.3. Refer to the NEWS file for further information:

http://icedtea.classpath.org/hg/release/icedtea7-2.3/file/icedtea-2.3.3/NEWS

All users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.125.4. RHBA-2012:1570 – java-1.7.0-openjdk bug fix update

Updated java-1.7.0-openjdk packages that fix one bug 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 Java Software Development Kit.

Bug Fix

BZ#880352

Previously, the Krb5LoginModule config class did not return a proper KDC list when krb5.conf file contained the "dns_lookup_kdc = true" property setting. With this update, a correct KDC list is returned under these circumstances.

All users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which fix this bug.

5.125.5. RHSA-2013:0165 – Important: java-1.7.0-openjdk security update

Updated java-1.7.0-openjdk packages that fix two security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

These packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Software Development Kit.

Security Fix

CVE-2012-3174, CVE-2013-0422

Two improper permission check issues were discovered in the reflection API in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

This erratum also upgrades the OpenJDK package to IcedTea7 2.3.4. Refer to the NEWS file, linked to in the References, for further information.

All users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.125.6. RHSA-2012:1223 – Important: java-1.7.0-openjdk security update

Updated java-1.7.0-openjdk packages that fix several security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

These packages provide the OpenJDK 7 Java Runtime Environment and the OpenJDK 7 Software Development Kit.

Security Fixes

CVE-2012-4681, CVE-2012-1682, CVE-2012-3136

Multiple improper permission check issues were discovered in the Beans component in OpenJDK. An untrusted Java application or applet could use these flaws to bypass Java sandbox restrictions.

CVE-2012-0547

A hardening fix was applied to the AWT component in OpenJDK, removing functionality from the restricted SunToolkit class that was used in combination with other flaws to bypass Java sandbox restrictions.

All users of java-1.7.0-openjdk are advised to upgrade to these updated packages, which resolve these issues. All running instances of OpenJDK Java must be restarted for the update to take effect.

5.126. JAVA-1.7.0-ORACLE

5.126.1. RHSA-2013:0237 – Critical: java-1.7.0-oracle security update

Updated java-1.7.0-oracle packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

Oracle Java SE version 7 includes the Oracle Java Runtime Environment and the Oracle Java Software Development Kit.

Security Fix

CVE-2012-1541, CVE-2012-3213, CVE-2012-3342, CVE-2013-0351, CVE-2013-0409, CVE-2013-0419, CVE-2013-0423, CVE-2013-0424, CVE-2013-0425, CVE-2013-0426, CVE-2013-0427, CVE-2013-0428, CVE-2013-0429, CVE-2013-0430, CVE-2013-0431, CVE-2013-0432, CVE-2013-0433, CVE-2013-0434, CVE-2013-0435, CVE-2013-0437, CVE-2013-0438, CVE-2013-0440, CVE-2013-0441, CVE-2013-0442, CVE-2013-0443, CVE-2013-0444, CVE-2013-0445, CVE-2013-0446, CVE-2013-0448, CVE-2013-0449, CVE-2013-0450, CVE-2013-1473, CVE-2013-1475, CVE-2013-1476, CVE-2013-1478, CVE-2013-1479, CVE-2013-1480, CVE-2013-1489

This update fixes several vulnerabilities in the Oracle Java Runtime Environment and the Oracle Java Software Development Kit. Further information about these flaws can be found on the Oracle Java SE Critical Patch Update Advisory page.

All users of java-1.7.0-oracle are advised to upgrade to these updated packages, which provide Oracle Java 7 Update 13 and resolve these issues. All running instances of Oracle Java must be restarted for the update to take effect.

5.126.2. RHSA-2012:1225 – Critical: java-1.7.0-oracle security update

Updated java-1.7.0-oracle packages that fix several security issues are now available for Red Hat Enterprise Linux 6 Supplementary.

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

The Oracle Java 7 release includes the Oracle Java 7 Runtime Environment and the Oracle Java 7 Software Development Kit.

Security Fix

CVE-2012-4681, CVE-2012-1682, CVE-2012-3136, CVE-2012-0547

This update fixes several vulnerabilities in the Oracle Java 7 Runtime Environment and the Oracle Java 7 Software Development Kit. Further information about these flaws can be found on the Oracle Java SE Security Alert page.

Red Hat is aware that a public exploit for CVE-2012-4681 is available that executes code without user interaction when a user visits a malicious web page using a browser with the Oracle Java 7 web browser plug-in enabled.

All users of java-1.7.0-oracle are advised to upgrade to these updated packages, which provide Oracle Java 7 Update 7 and resolve these issues. All running instances of Oracle Java must be restarted for the update to take effect.

5.126.3. RHSA-2012:1391 – Critical: java-1.7.0-oracle security update

Updated java-1.7.0-oracle packages that fix several security issues are now available for Red Hat Enterprise Linux 6 Supplementary.

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

Oracle Java SE version 7 includes the Oracle Java Runtime Environment and the Oracle Java Software Development Kit.

Security Fix

CVE-2012-1531, CVE-2012-1532, CVE-2012-1533, CVE-2012-3143, CVE-2012-3159, CVE-2012-3216, CVE-2012-4416, CVE-2012-5067, CVE-2012-5068, CVE-2012-5069, CVE-2012-5070, CVE-2012-5071, CVE-2012-5072, CVE-2012-5073, CVE-2012-5074, CVE-2012-5075, CVE-2012-5076, CVE-2012-5077, CVE-2012-5079, CVE-2012-5081, CVE-2012-5083, CVE-2012-5084, CVE-2012-5085, CVE-2012-5086, CVE-2012-5087, CVE-2012-5088, CVE-2012-5089

This update fixes several vulnerabilities in the Oracle Java Runtime Environment and the Oracle Java Software Development Kit. Further information about these flaws can be found on the Oracle Java SE Critical Patch Update Advisory page.

All users of java-1.7.0-oracle are advised to upgrade to these updated packages, which provide Oracle Java 7 Update 9. All running instances of Oracle Java must be restarted for the update to take effect.

5.126.4. RHSA-2013:0156 – Critical: java-1.7.0-oracle security update

Updated java-1.7.0-oracle packages that fix two security issues are now available for Red Hat Enterprise Linux 5 and 6 Supplementary.

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

Oracle Java SE version 7 includes the Oracle Java Runtime Environment and the Oracle Java Software Development Kit.

Security Fix

CVE-2012-3174, CVE-2013-0422

This update fixes two vulnerabilities in the Oracle Java Runtime Environment and the Oracle Java Software Development Kit. Further information about these flaws can be found on the Oracle Security Alert page.

Red Hat is aware that a public exploit for CVE-2013-0422 is available that executes code without user interaction when a user visits a malicious web page using a browser with the Oracle Java 7 web browser plug-in enabled.

All users of java-1.7.0-oracle are advised to upgrade to these updated packages, which provide Oracle Java 7 Update 11 and resolve these issues. All running instances of Oracle Java must be restarted for the update to take effect.

5.127. JSS

5.127.1. RHBA-2012:0920 – jss bug fix update

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

JSS (Java Security Services) is a Java binding to Network Security Services (NSS), which provides SSL/TLS network protocols and other security services in the Public Key Infrastructure (PKI) suite. JSS is primarily utilized by the Certificate Server.

Bug Fixes

BZ#767768

During key archival process, DRM (Data Recovery Manager) decrypted user's private keys and then re-encrypted the keys for storage purposes. The reverse process took place during key recovery; therefore, the private key was not processed in a token at all times as the decrypted private key was present in the DRM memory between the time of decryption and encryption. This update adds the secure PKCS #12 and PKCS #5 v2.0 support, support for wrapping and unwrapping private keys in their token, and secure private key handling for TMS (Token Management System) key recovery to Red Hat Certificate System 8.1. As a result, the key archival operations now happen in the token.

BZ#767771

The "kra.storageUnit.hardware" configuration parameter did not exist in DRM's CS.cfg after upgrade. Consequently, if parameter "kra.storageUnit.hardware" was defined, recovery operations failed and the server returned the following error message:

> PKCS #12 Creation Failed java.lang.IllegalArgumentException: bagType or bagContent is null

This update modifies the jss, pki-kra, pki-common components so that the "kra.storageUnit.hardware" configuration parameter is processed correctly. As a result, the key archival and recovery process is successful on in-place upgraded and migrated instances.

BZ#767773

Previously, JSS was using the HSM (Hardware Security Module) token name as manufacturer ID. If the HSM token name differed from the manufacturer ID, the key archival and recovery failed. This update adds logic to JSS so that it can recognize the currently supported HSMs: nCipher and SafeNet. Key archival and recovery in TMS and non-TMS Common Criteria environments now work as expected.

All users of jss are advised to upgrade to these updated packages, which fix these bugs.

5.128. KABI-WHITELISTS

5.128.1. RHEA-2012:0918 – kabi-whitelists enhancement update

An updated kabi-whitelists package that adds various enhancements is now available for Red Hat Enterprise Linux 6.

The kabi-whitelists package contains reference files documenting interfaces provided by the Red Hat Enterprise Linux 6 kernel that are considered to be stable by Red Hat engineering, and safe for longer term use by third-party loadable device drivers, as well as for other purposes.

Enhancements

BZ#722619

Multiple symbols have been added to the Red Hat Enterprise Linux 6.3 kernel application binary interface (ABI) whitelists.

BZ#737276

Multiple symbols for Hitachi loadable device drivers have been added to the kernel ABI whitelists.

BZ#753771

This update modifies the structure of the kabi-whitelists package: whitelists are now ordered according to various Red Hat Enterprise Linux releases, and a symbolic link that points to the latest release has been added.

BZ#803885

The "__dec_zone_page_state" and "dec_zone_page_state" symbols have been added to the kernel ABI whitelists.

BZ#810456

The "blk_queue_rq_timed_out", "fc_attach_transport", "fc_release_transport", "fc_remote_port_add", "fc_remote_port_delete", "fc_remote_port_rolechg", "fc_remove_host", and "touch_nmi_watchdog" symbols have been added to the kernel ABI whitelists.

BZ#812463

Multiple symbols for Oracle Cloud File System have been added to the kernel ABI whitelists.

BZ#816533

The "get_fs_type" and "vscnprintf" have been added to the kernel ABI whitelists.

All users of kabi-whitelists are advised to upgrade to this updated package, which adds these enhancements.

5.129. KDEARTWORK

5.129.1. RHBA-2012:0450 – kdeartwork bug fix update

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

The K Desktop Environment (KDE) is a graphical desktop environment for the X Window System. The kdeartwork packages include styles, themes and screen savers for KDE.

Bug Fix

BZ#736624

Previously, the KPendulum and KRotation screen savers, listed in the OpenGL group of KDE screen savers, produced only a blank screen. This update disables KPendulum and KRotation and none of them is listed in the OpenGL group anymore.

All users of kdeartwork are advised to upgrade to these updated packages, which fix this bug.

5.130. KDEBASE

5.130.1. RHBA-2012:1371 – kdebase bug fix update

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

The K Desktop Environment (KDE) is a graphical desktop environment for the X Window System. The kdebase packages include core applications for KDE.

Bug Fixes

BZ#608007

Prior to this update, the Konsole context menu item "Show menu bar" was always checked in new windows even if this menu item was disabled before. This update modifies the underlying code to handle the menu item "Show menu bar" as expected.

BZ#729307

Prior to this update, users could not define a default size for xterm windows when using the Konsole terminal in KDE. This update modifies the underlying code and adds the functionality to define a default size.

All users of kdebase are advised to upgrade to these updated packages, which fix these bugs.

5.131. KDEBASE-WORKSPACE

5.131.1. RHBA-2012:1286 – kdebase-workspace bug fix update

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

The kdebase-workspace packages contain utilities for basic operations with the desktop environment. The utilities allow users for example, to change system settings, resize and rotate X screens or set panels and widgets on the workspace.

Bug Fix

BZ#749460

Prior to this update, the task manager did not honor the order of manually arranged items. As a consequence, manually arranged taskbar entries were randomly rearranged when the user switched desktops. This update modifies the underlying code to make manually arranged items more persistent.

All users of kdebase-workspace are advised to upgrade to these updated packages, which fix this bug.

5.131.2. RHBA-2012:0400 – kdebase-workspace bug fix update

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

KDE is a graphical desktop environment for the X Window System. The kdebase-workspace packages contain utilities for basic operations with the desktop environment. The utilities allow users for example, to change system settings, resize and rotate X screens or set panels and widgets on the workspace.

Bug Fix

BZ#724960

Previously, the kdebase-workspace package relied on the bluez-libs-devel package for rebuild. However, bluez-libs-devel was not supported on IBM System z architectures and builds could be created only with help of the fake-build-provides package which is not required behavior. With this update, the bluez-libs-devel package is no longer required as a dependency on IBM System z architecture and rebuilds are successful.

All users of kdebase-workspace are advised to upgrade to these updated packages, which fix this bug.

5.132. KDELIBS3

5.132.1. RHBA-2012:1244 – kdelibs3 bug fix update

Updated kdelibs3 packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The kdelibs3 packages provide libraries for the K Desktop Environment (KDE).

Bug Fixes

BZ#681901

Prior to this update, the kdelibs3 libraries caused a conflict for the subversion version control tool. As a consequence, subvervision was not correctly built if the kdelibs3 libraries were installed. This update modifies the underlying code to avoid this conflict. Now, subversion builds as expected with kdelibs3.

BZ#734447

kdelibs3 provided its own set of trusted Certificate Authority (CA) certificates. This update makes kdelibs3 use the system set from the ca-certificates package, instead of its own copy.

All users of kdelibs3 are advised to upgrade to these updated packages, which fix these bugs.

5.133. KDELIBS

5.133.1. RHSA-2012:1416 – Critical: kdelibs security update

Updated kdelibs packages that fix two security issues are now available for Red Hat Enterprise Linux 6.

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

The kdelibs packages provide libraries for the K Desktop Environment (KDE). Konqueror is a web browser.

Security Fixes

CVE-2012-4512

A heap-based buffer overflow flaw was found in the way the CSS (Cascading Style Sheets) parser in kdelibs parsed the location of the source for font faces. A web page containing malicious content could cause an application using kdelibs (such as Konqueror) to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2012-4513

A heap-based buffer over-read flaw was found in the way kdelibs calculated canvas dimensions for large images. A web page containing malicious content could cause an application using kdelibs to crash or disclose portions of its memory.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The desktop must be restarted (log out, then log back in) for this update to take effect.

5.133.2. RHBA-2012:1251 – kdelibs bug fix update

Updated kdelibs packages that fix various bugs are now available for Red Hat Enterprise Linux 6.

The kdelibs packages provide libraries for the K Desktop Environment (KDE).

Bug Fixes

BZ#587016

Prior to this update, the KDE Print dialog did not remember previous settings, nor did it allow the user

to save the settings. Consequent to this, when printing several documents, users were forced to manually change settings for each printed document. With this update, the KDE Print dialog retains previous settings as expected.

BZ#682611

When the system was configured to use the Traditional Chinese language (the zh_TW locale), Konqueror incorrectly used a Chinese (zh_CN) version of its splash page. This update ensures that Konqueror uses the correct locale.

BZ#734734

Previously, clicking the system tray to display hidden icons could cause the Plasma Workspaces to consume an excessive amount of CPU time. This update applies a patch that fixes this error.

BZ#754161

When using Konqueror to recursively copy files and directories, if one of the subdirectories was not accessible, no warning or error message was reported to the user. This update ensures that Konqueror displays a proper warning message in this scenario.

BZ#826114

Prior to this update, an attempt to add "Terminal Emulator" to the Main Toolbar caused Konqueror to terminate unexpectedly with a segmentation fault. With this update, the underlying source code has been corrected to prevent this error so that users can now use this functionality as expected.

All users of kdelibs are advised to upgrade to these updated packages, which fix these bugs.

5.133.3. RHSA-2012:1418 – Critical: kdelibs security update

Updated kdelibs packages that fix two security issues are now available for Red Hat Enterprise Linux 6 FasTrack.

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

The kdelibs packages provide libraries for the K Desktop Environment (KDE). Konqueror is a web browser.

Security Fixes

CVE-2012-4512

A heap-based buffer overflow flaw was found in the way the CSS (Cascading Style Sheets) parser in kdelibs parsed the location of the source for font faces. A web page containing malicious content could cause an application using kdelibs (such as Konqueror) to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2012-4513

A heap-based buffer over-read flaw was found in the way kdelibs calculated canvas dimensions for large images. A web page containing malicious content could cause an application using kdelibs to crash or disclose portions of its memory.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The desktop must be restarted (log out, then log back in) for this update to take effect.

5.133.4. RHBA-2012:0377 – kdelibs bug fix update

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

The kdelibs packages provide libraries for K Desktop Environment (KDE).

Bug Fix

BZ#698286

Previously, on big-endian architectures, including IBM System z, the Konqueror web browser could terminate unexpectedly or become unresponsive when loading certain web sites. A patch has been applied to address this issue, and Konqueror no longer crashes or hangs on the aforementioned architectures.

All users of kdelibs are advised to upgrade to these updated packages, which fix this bug.

5.134. KDEPIM

5.134.1. RHBA-2012:1287 – kdepim bug fix update

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

The KDE Personal Information Management (kdepim) suite helps to organize your mail, tasks, appointments, and contacts.

Bug Fix

BZ#811125

Prior to this update, the cyrus-sasl-plain package was not a dependency of the kdepim package. As a consequence, Kmail failed to send mail. This update modifies the underlying code to include the cyrus-sasl-plain dependency.

All users of kdepim are advised to upgrade to these updated packages, which fix this bug.

5.135. KERNEL

5.135.1. RHSA-2013-1783 – 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 6.3 Extended Update Support.

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-2012-4508, Important

A race condition was found in the way asynchronous I/O and fallocate() interacted when using the ext4 file system. A local, unprivileged user could use this flaw to expose random data from an extent whose data blocks have not yet been written, and thus contain data from a deleted file.

CVE-2013-4299, Moderate

An information leak flaw was found in the way the Linux kernel's device mapper subsystem, under certain conditions, interpreted data written to snapshot block devices. An attacker could use this flaw to read data from disk blocks in free space, which are normally inaccessible.

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 Theodore Ts'o for reporting CVE-2012-4508, Fujitsu for reporting CVE-2013-4299, and Kees Cook for reporting CVE-2013-2851. Upstream acknowledges Dmitry Monakhov as the original reporter of CVE-2012-4508.

Bug Fixes

BZ#1016105

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 running 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#1017505, BZ#1017506

A previous change in the port auto-selection code allowed sharing of 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#1017903

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 could 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#1020527

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#1023349

Previously, when the user added an IPv6 route for local delivery, the route did not work and packets could not be sent. A patch has been applied to limit the neighbor entry creation only for input flow, thus fixing this bug. As a result, IPv6 routes for local delivery now work as expected.

BZ#1028592

A bug in the kernel's file system code allowed the d_splice_alias() function to create a new dentry for a directory with an already-existing non-DISCONNECTED dentry. As a consequence, a thread accessing the directory could attempt to take the i_mutex on that directory twice, resulting in a deadlock situation. To resolve this problem, d_splice_alias() has been modified so that in the problematic cases, it reuses an existing dentry instead of creating a new dentry.

BZ#1029423

The kernel's thread helper previously used larvals of the request threads without holding a reference count. This could result in a NULL pointer dereference and subsequent kernel panic if the helper thread completed after the larval had been destroyed upon the request thread exiting. With this update, the helper thread holds a reference count on the request threads larvals so that a NULL pointer dereference is now avoided.

BZ#1029901

Due to a bug in the SELinux Makefile, a kernel compilation could fail when the "-j" option was specified to perform the compilation with multiple parallel jobs. This happened because SELinux expected the existence of an automatically generated file, "flask.h", prior to the compiling of some dependent files. The Makefile has been corrected and the "flask.h" dependency now applies to all objects from the "selinux-y" list. The parallel compilation of the kernel now succeeds as expected.

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.

5.135.2. RHBA-2013:1104 – kernel bug fix update

Updated kernel packages that fix several bugs are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Bug Fixes

BZ#969341

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#972599

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#975577

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#976695

* The schedule_ipi() function is called in the hardware interrupt context and it raises the SCHED_SOFTIRQ software interrupts to perform system load balancing. Software interrupts in Linux are either performed on return from a hardware interrupt or are handled by the ksoftirqd daemon if the interrupts cannot be processed normally. Previously, the context of the schedule_ipi() function was not marked as a hardware interrupt so while performing schedule_ipi(), the ksoftirqd daemon could have been triggered. When triggered, the daemon attempted to balance the system load. However at that time, the load balancing had already been performed by the SCHED_SOFTIRQ software interrupt so the ksoftirqd daemon attempted to balance the already-balanced system, which led to excessive consumption of CPU time. The problem has been resolved by adding irq_enter() and irq_exit() function calls to schedule IPI handlers, which assures that context of softirq_ipi() is correctly marked as a hardware interrupt and the ksoftirqd daemon is no longer triggered when the SCHED_SOFTIRQ interrupt has been raised.

BZ#977667

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.

Users should upgrade to these updated packages, which contain backported patches to correct these bugs. The system must be rebooted for this update to take effect.

5.135.3. RHSA-2013:0928 – Important: kernel security and bug fix update

Updated kernel packages that fix several security issues and bugs are now available for Red Hat Enterprise Linux 6 Extended Update Support.

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-2013-0311, Important

A flaw was found in the way the vhost kernel module handled descriptors that spanned multiple regions. A privileged guest user in a KVM (Kernel-based Virtual Machine) guest could use this flaw to crash the host or, potentially, escalate their privileges on the host.

CVE-2013-1773, Important

A buffer overflow flaw was found in the way UTF-8 characters were converted to UTF-16 in the utf8s_to_utf16s() function of the Linux kernel's FAT file system implementation. A local user able to mount a FAT file system with the "utf8=1" option could use this flaw to crash the system or, potentially, to escalate their privileges.

CVE-2013-1796, Important

A flaw was found in the way KVM handled guest time updates when the buffer the guest registered by writing to the MSR_KVM_SYSTEM_TIME machine state register (MSR) crossed a page boundary. A privileged guest user could use this flaw to crash the host or, potentially, escalate their privileges, allowing them to execute arbitrary code at the host kernel level.

CVE-2013-1797, Important

A potential use-after-free flaw was found in the way KVM handled guest time updates when the GPA (guest physical address) the guest registered by writing to the MSR_KVM_SYSTEM_TIME machine state register (MSR) fell into a movable or removable memory region of the hosting user-space process (by default, QEMU-KVM) on the host. If that memory region is deregistered from KVM using KVM_SET_USER_MEMORY_REGION and the allocated virtual memory reused, a privileged guest user could potentially use this flaw to escalate their privileges on the host.

CVE-2013-1798, Important

A flaw was found in the way KVM emulated IOAPIC (I/O Advanced Programmable Interrupt Controller). A missing validation check in the ioapic_read_indirect() function could allow a privileged guest user to crash the host, or read a substantial portion of host kernel memory.

CVE-2012-4542, Moderate

It was found that the default SCSI command filter does not accommodate commands that overlap across device classes. A privileged guest user could potentially use this flaw to write arbitrary data to a LUN that is passed-through as read-only.

CVE-2013-1767, Low

A use-after-free flaw was found in the tmpfs implementation. A local user able to mount and unmount a tmpfs file system could use this flaw to cause a denial of service or, potentially, escalate their privileges.

CVE-2013-1848, Low

A format string flaw was found in the ext3_msg() function in the Linux kernel's ext3 file system implementation. A local user who is able to mount an ext3 file system could use this flaw to cause a denial of service or, potentially, escalate their privileges.

Red Hat would like to thank Andrew Honig of Google for reporting the CVE-2013-1796, CVE-2013-1797, and CVE-2013-1798 issues. The CVE-2012-4542 issue was discovered by Paolo Bonzini of Red Hat.

Bug Fixes

BZ#952612

When pNFS (parallel NFS) code was in use, a file locking process could enter a deadlock while trying to recover form a server reboot. This update introduces a new locking mechanism that avoids the deadlock situation in this scenario.

BZ#955503

The be2iscsi driver previously leaked memory in the driver's control path when mapping tasks.This update fixes the memory leak by freeing all resources related to a task when the task was completed. Also, the driver did not release a task after responding to the received NOP-IN acknowledgment with a valid Target Transfer Tag (TTT). Consequently, the driver run out of tasks available for the session and no more iscsi commands could be issued. A patch has been applied to fix this problem by releasing the task.

BZ#956295

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#956933

A bug in the lpfc driver allowed re-enabling of an interrupt from the interrupt context so the interrupt handler was able to re-enter the interrupt context. The interrupt context re-entrance problem led to kernel stack corruption which consequently resulted in a kernel panic. This update provides a patch addressing the re-entrance problem so the kernel stack corruption and the subsequent kernel panic can no longer occur under these circumstances.

BZ#960410

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#960416

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#960419

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 caused by incorrect error handling in the page cache. This problem has been fixed and the correct error value is now returned.

BZ#960424

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#962367

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#964960

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.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

5.135.4. RHBA-2013:0768 – kernel bug fix update

Updated kernel packages that fix several bugs are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Bug Fixes

BZ#911266

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 to turn off the feature and update the BIOS.

BZ#920264

The NFS code implements the "silly rename" operation to handle an open file that is held by a

process while another process attempts to remove it. The "silly rename" operation works according to the "delete on last close" semantics so the inode of the file is not released until the last process that opens the file closes it. A previous update of the NFS code broke the mechanics that prevented an NFS client from deleting a silly-renamed entry. This affected the "delete on last close" semantics and silly-renamed files could be deleted by any process while the files were open for I/O by another process. As a consequence, the process reading the file failed with the "ESTALE" error code. This update modifies the way the NFS code handles dentries of silly-renamed files and silly-renamed files can not be deleted until the last process that has the file open for I/O closes it.

BZ#920267

The NFSv4 code uses byte range locks to simulate the flock() function, which is used to apply or remove an exclusive advisory lock on an open file. However, using the NFSv4 byte range locks precludes a possibility to open a file with read-only permissions and subsequently to apply an exclusive advisory lock on the file. A previous patch broke a mechanism used to verify the mode of the open file. As a consequence, the system became unresponsive and the system logs filled with a "kernel: nfs4_reclaim_open_state: Lock reclaim failed!" error message if the file was open with read-only permissions and an attempt to apply an exclusive advisory lock was made. This update modifies the NFSv4 code to check the mode of the open file before attempting to apply the exclusive advisory lock. The "-EBADF" error code is returned if the type of the lock does not match the file mode.

BZ#921960

When running a high thread workload of small-sized files on an XFS file system, 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 the kernel no longer panics in this scenario.

BZ#923850

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 resends blocking lock requests after the grace period has expired.

BZ#924838

A bug in the anon_vma lock in the mprotect() function could cause virtual memory area (vma) corruption. The bug has been fixed so that virtual memory area corruption no longer occurs in this scenario.

All users are advised to upgrade to these updated packages, which fix these bugs. The system must be rebooted for this update to take effect.

5.135.5. RHSA-2012:1366 – Important: kernel security and bug fix update

Updated kernel 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 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 kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2012-3412, Important

A flaw was found in the way socket buffers (skb) requiring TSO (TCP segment offloading) were handled by the sfc driver. If the skb did not fit within the minimum-size of the transmission queue, the network card could repeatedly reset itself. A remote attacker could use this flaw to cause a denial of service.

Red Hat would like to thank Ben Hutchings of Solarflare for reporting this issue.

Bug Fixes

BZ#856316

In Fibre Channel fabrics with large zones, the automatic port rescan on incoming Extended Link Service (ELS) frames and any adapter recovery could cause high traffic, in particular if many Linux instances shared a host bus adapter (HBA), which is common on IBM System z architecture. This could lead to various failures; for example, names server requests, port or adapter recovery could fail. With this update, ports are re-scanned only when setting an adapter online or on manual user-triggered writes to the sysfs attribute "port_rescan".

BZ#856686

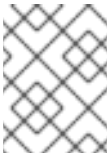
Under certain circumstances, a system crash could result in data loss on XFS file systems. If files were created immediately before the file system was left to idle for a long period of time and then the system crashed, those files could appear as zero-length once the file system was remounted. This occurred even if a sync or fsync was run on the files. This was because XFS was not correctly idling the journal, and therefore it incorrectly replayed the inode allocation transactions upon mounting after the system crash, which zeroed the file size. This problem has been fixed by re-instating the periodic journal idling logic to ensure that all metadata is flushed within 30 seconds of modification, and the journal is updated to prevent incorrect recovery operations from occurring.

BZ#856703

On architectures with the 64-bit cputime_t type, it was possible to trigger the "divide by zero" error, namely, on long-lived processes. A patch has been applied to address this problem, and the "divide by zero" error no longer occurs under these circumstances.

BZ#857012

The kernel provided by the Red Hat Enterprise Linux 6.3 release included an unintentional kernel ABI (kABI) breakage with regards to the "contig_page_data" symbol. Unfortunately, this breakage did not cause the checksums to change. As a result, drivers using this symbol could silently corrupt memory on the kernel. This update reverts the previous behavior.



NOTE

Any driver compiled with the "contig_page_data" symbol during the early release of Red Hat Enterprise Linux 6.3 needs to be recompiled again for this kernel.

BZ#857334

A race condition could occur between page table sharing and virtual memory area (VMA) teardown. As a consequence, multiple "bad pmd" message warnings were displayed and "kernel BUG at mm/filemap.c:129" was reported while shutting down applications that share memory segments backed by huge pages. With this update, the VM_MAYSHARE macro is explicitly cleaned during the unmap_hugepage_range() call under the i_mmap_lock. This makes VMA ineligible for sharing and avoids the race condition. After using shared segments backed by huge pages, applications like databases and caches shut down correctly, with no crash.

BZ#857854

A kernel panic could occur when using the be2net driver. This was because the Bottom Half (BF) was enabled even if the Interrupt ReQuest (IRQ) was already disabled. With this update, the BF is disabled in callers of the be_process_mcc() function and the kernel no longer crashes in this scenario.



NOTE

Note that, in certain cases, it is possible to experience the network card being unresponsive after installing this update. A future update will correct this problem.

BZ#858284

The Stream Control Transmission Protocol (SCTP) ipv6 source address selection logic did not take the preferred source address into consideration. With this update, the source address is chosen from the routing table by taking this aspect into consideration. This brings the SCTP source address selection on par with IPv4.

BZ#858285

Prior to this update, it was not possible to set IPv6 source addresses in routes as it was possible with IPv4. With this update, users can select the preferred source address for a specific IPv6 route with the "src" option of the "ip -6 route" command.

All users of kernel should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

5.135.6. RHSA-2012:1304 – Moderate: 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 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 kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2012-2313, Low

A flaw was found in the way the Linux kernel's dl2k driver, used by certain D-Link Gigabit Ethernet adapters, restricted IOCTLs. A local, unprivileged user could use this flaw to issue potentially harmful IOCTLs, which could cause Ethernet adapters using the dl2k driver to malfunction (for example, losing network connectivity).

CVE-2012-2384, Moderate

An integer overflow flaw was found in the i915_gem_do_execbuffer() function in the Intel i915 driver in the Linux kernel. A local, unprivileged user could use this flaw to cause a denial of service. This issue only affected 32-bit systems.

CVE-2012-2390, Moderate

A memory leak flaw was found in the way the Linux kernel's memory subsystem handled resource clean up in the mmap() failure path when the MAP_HUGETLB flag was set. A local, unprivileged user could use this flaw to cause a denial of service.

CVE-2012-3430, Low

A flaw was found in the way the msg_namelen variable in the rds_recvmsg() function of the Linux kernel's Reliable Datagram Sockets (RDS) protocol implementation was initialized. A local, unprivileged user could use this flaw to leak kernel stack memory to user-space.

CVE-2012-3552, Moderate

A race condition was found in the way access to inet->opt ip_options was synchronized in the Linux kernel's TCP/IP protocol suite implementation. Depending on the network facing applications running on the system, a remote attacker could possibly trigger this flaw to cause a denial of service. A local, unprivileged user could use this flaw to cause a denial of service regardless of the applications the system runs.

Red Hat would like to thank Hafid Lin for reporting CVE-2012-3552, and Stephan Mueller for reporting CVE-2012-2313. The CVE-2012-3430 issue was discovered by the Red Hat InfiniBand team.

Bug Fixes

BZ#812962

Previously, after a crash, preparing to switch to the kdump kernel could in rare cases race with IRQ migration, causing a deadlock of the ioapic_lock variable. As a consequence, kdump became unresponsive. The race condition has been fixed, and switching to kdump no longer causes hangs in this scenario.

BZ#842757

The xmit packet size was previously 64K, exceeding the hardware capability of the be2net card because the size did not account for the Ethernet header. The adapter was therefore unable to process xmit requests exceeding this size, produced error messages and could become unresponsive. To prevent these problems, GSO (Generic Segmentation Offload) maximum size has been reduced to account for the Ethernet header.

BZ#842982

When the netconsole module was configured over bridge and the "service network restart" command was executed, a deadlock could occur, resulting in a kernel panic. This was caused by recursive rtnl locking by both bridge and netconsole code during network interface unregistration. With this update, the rtnl lock usage is fixed, and the kernel no longer crashes in this scenario.

BZ#842984

When using virtualization with the netconsole module configured over the main system bridge, guests could not be added to the bridge, because TAP interfaces did not support netpoll. This update adds support of netpoll to the TUN/TAP interfaces so that bridge devices in virtualization

setups can use netconsole.

BZ#843102

Signed-unsigned values comparison could under certain circumstances lead to a superfluous reshed_task() routine to be called, causing several unnecessary cycles in the scheduler. This problem has been fixed, preventing the unnecessary cycles in the scheduler.

BZ#845464

If RAID1 or RAID10 was used under LVM or some other stacking block device, it was possible to enter a deadlock during a resync or recovery operation. Consequently, md RAID devices could become unresponsive on certain workloads. This update avoids the deadlock so that md RAID devices work as expected under these circumstances.

BZ#846216

Previously, soft interrupt requests (IRQs) under the bond_alb_xmit() function were locked even when the function contained soft IRQs that were disabled. This could cause a system to become unresponsive or terminate unexpectedly. With this update, such IRQs are no longer disabled, and the system no longer hangs or crashes in this scenario.

BZ#846832

Previously, the TCP socket bound to NFS server contained a stale skb_hints socket buffer. Consequently, kernel could terminate unexpectedly. A patch has been provided to address this issue and skb_hints is now properly cleared from the socket, thus preventing this bug.

BZ#846836

A race condition could occur due to incorrect locking scheme in the code for software RAID. Consequently, this could cause the mkfs utility to become unresponsive when creating an ext4 file system on software RAID5. This update introduces a locking scheme in the handle_stripe() function, which ensures that the race condition no longer occurs.

BZ#846838

When a device is added to the system at runtime, the AMD IOMMU driver initializes the necessary data structures to handle translation for it. Previously, however, the per-device dma_ops structure types were not changed to point to the AMD IOMMU driver, so mapping was not performed and direct memory access (DMA) ended with the IO_PAGE_FAULT message. This consequently led to networking problems. With this update, the structure types point correctly to the AMD IOMMU driver, and networking works as expected when the AMD IOMMU driver is used.

BZ#846839

Due to an error in the dm-mirror driver, when using LVM mirrors on disks with discard support (typically SSD disks), repairing such disks caused the system to terminate unexpectedly. The error in the driver has been fixed and repairing disks with discard support is now successful.

BZ#847042

On Intel systems with Pause Loop Exiting (PLE), or AMD systems with Pause Filtering (PF), it was possible for larger multi-CPU KVM guests to experience slowdowns and soft lock-ups. Due to a boundary condition in kvm_vcpu_on_spin, all the VCPUs could try to yield to VCPU0, causing contention on the run queue lock of the physical CPU where the guest's VCPU0 is running. This update eliminates the boundary condition in kvm_vcpu_on_spin.

BZ#847045

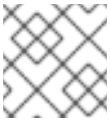
Previously, using the e1000e driver could lead to a kernel panic. This was caused by a NULL pointer dereference that occurred if the adapter was being closed and reset simultaneously. The source code of the driver has been modified to address this problem, and kernel no longer crashes in this scenario.

BZ#847727

On PowerPC architecture, the "top" utility displayed incorrect values for the CPU idle time, delays and workload. This was caused by a previous update that used jiffies for the I/O wait and idle time, but the change did not take into account that jiffies and CPU time are represented by different units. These differences are now taken into account, and the "top" utility displays correct values on PowerPC architecture.

BZ#847945

Due to a missing return statement, the nfs_attr_use_mounted_on_file() function returned a wrong value. As a consequence, redundant ESTALE errors could potentially be returned. This update adds the proper return statement to nfs_attr_use_mounted_on_file(), thus preventing this bug.



NOTE

This bug only affected NFS version 4 file systems.

BZ#849051

A deadlock sometimes occurred between the dlm_controld daemon closing a lowcomms connection through the configfs file system and the dlm_send process looking up the address for a new connection in configfs. With this update, the node addresses are saved within the lowcomms code so that the lowcomms work queue does not need to use configfs to get a node address.

BZ#849551

Performance of O_DSYNC on the GFS2 file system was affected when only data (not metadata such as file size) was dirtied as a result of a write system call. This was because O_DSYNC writes were always behaving in the same way as O_SYNC. With this update, O_DSYNC writes only write back data, if the inode's metadata is not dirty. This leads to a considerable performance improvement in this case. Note that this problem does not affect data integrity. The same issue also applies to the pairing of write and fdatasync calls.

BZ#851444

If a mirror or redirection action is configured to cause packets to go to another device, the classifier holds a reference count. However, it was previously assuming that the administrator cleaned up all redirections before removing. Packets were therefore dropped if the mirrored device was not present, and connectivity to the host could be lost. To prevent such problems, a notifier and cleanup are now run during the unregister action. Packets are not dropped if the a mirrored device is not present.

BZ#851445

The kernel contains a rule to blacklist direct memory access (DMA) modes for "2GB ATA Flash Disk" devices. However, this device ID string did not contain a space at the beginning of the name. Due to this, the rule failed to match the device and failed to disable DMA modes. With this update, the string correctly reads " 2GB ATA Flash Disk", and the rule can be matched as expected.

All users of kernel should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

5.135.7. RHBA-2012:1104 – kernel bug fix update

Updated kernel packages that fix four bugs are now available for Red Hat Enterprise Linux 6.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Bug Fixes

BZ#836904

Previously, futex operations on read-only (RO) memory maps did not work correctly. This broke workloads that had one or more reader processes performing the FUTEX_WAIT operation on a futex within a read-only shared file mapping and a writer process that had a writable mapping performing the FUTEX_WAKE operation. With this update, the FUTEX_WAKE operation is performed with a RO MAP_PRIVATE mapping, and is successfully awaken if another process updates the region of the underlying mapped file.

BZ#837218

When removing a bonding module, the bonding driver uses code separate from the net device operations to clean up the VLAN code. Recent changes to the kernel introduced a bug which caused a kernel panic if the vlan module was removed after the bonding module had been removed. To fix this problem, the VLAN group removal operations found in the bonding kill_vid path are now duplicated in alternate paths which are used when removing a bonding module.

BZ#837227

The bonding method for adding VLAN Identifiers (VIDs) did not always add the VID to a slave VLAN group. When the NETIF_F_HW_VLAN_FILTER flag was not set on a slave, the bonding module could not add new VIDs to it. This could cause networking problems and the system to be unreachable even if NIC messages did not indicate any problems. This update changes the bond VID add path to always add a new VID to the slaves (if the VID does not exist). This ensures that networking problems no longer occur in this scenario.

BZ#837843

Previously, reference counting was imbalanced in the slave add and remove paths for bonding. If a network interface controller (NIC) did not support the NETIF_F_HW_VLAN_FILTER flag, the bond_add_vlans_on_slave() and bond_del_vlans_on_slave() functions did not work properly, which could lead to a kernel panic if the VLAN module was removed while running. The underlying source code for adding and removing a slave and a VLAN has been revised and now also contains a common path, so that kernel crashes no kernel no longer occur in the described scenario.

All users of kernel are advised to upgrade to these updated packages, which fix these bugs. The system must be rebooted for this update to take effect.

5.135.8. RHSA-2013:0223 – Moderate: kernel security and bug fix update

Updated kernel packages that fix three security issues and multiple 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.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2012-4398, Moderate

It was found that a deadlock could occur in the Out of Memory (OOM) killer. A process could trigger this deadlock by consuming a large amount of memory, and then causing request_module() to be called. A local, unprivileged user could use this flaw to cause a denial of service (excessive memory consumption).

CVE-2012-4461, Moderate

A flaw was found in the way the KVM (Kernel-based Virtual Machine) subsystem handled guests attempting to run with the X86_CR4_OSXSAVE CPU feature flag set. On hosts without the XSAVE CPU feature, a local, unprivileged user could use this flaw to crash the host system. (The "grep --color xsave /proc/cpuinfo" command can be used to verify if your system has the XSAVE CPU feature.)

CVE-2012-4530, Low

A memory disclosure flaw was found in the way the load_script() function in the binfmt_script binary format handler handled excessive recursions. A local, unprivileged user could use this flaw to leak kernel stack memory to user-space by executing specially-crafted scripts.

Red Hat would like to thank Tetsuo Handa for reporting CVE-2012-4398, and Jon Howell for reporting CVE-2012-4461.

Bug Fixes

BZ#846840

When an NFSv4 client received a read delegation, a race between the OPEN and DELEGRETURN operation could occur. If the DELEGRETURN operation was processed first, the NFSv4 client treated the delegation returned by the following OPEN as a new delegation. Also, the NFSv4 client did not correctly handle errors caused by requests that used a bad or revoked delegation state ID. As a result, applications running on the client could receive spurious EIO errors. This update applies a series of patches that fix the NFSv4 code so an NFSv4 client recovers correctly in the described situations instead of returning errors to applications.

BZ#865305

Filesystem in Userspace (FUSE) did not implement scatter-gather direct I/O optimally. Consequently, the kernel had to process an extensive number of FUSE requests, which had a negative impact on system performance. This update applies a set of patches which improves internal request management for other features, such as readahead. FUSE direct I/O overhead has been significantly reduced to minimize negative effects on system performance.

BZ#876090

In case of a regular CPU hot plug event, the kernel does not keep the original cpuset configuration and can reallocate running tasks to active CPUs. Previously, the kernel treated switching between suspend and resume modes as a regular CPU hot plug event, which could have a significant negative impact on system performance in certain environments such as SMP KVM virtualization. When

resuming an SMP KVM guest from suspend mode, the libvirtd daemon and all its child processes were pinned to a single CPU (the boot CPU) so that all VMs used only the single CPU. This update applies a set of patches which ensure that the kernel does not modify cpusets during suspend and resume operations. The system is now resumed in the exact state before suspending without any performance decrease.

BZ#878774

Previously, the kernel had no way to distinguish between a device I/O failure due to a transport problem and a failure as a result of command timeout expiration. I/O errors always resulted in a device being set offline and the device had to be brought online manually even though the I/O failure occured due to a transport problem. With this update, the SCSI driver has been modified and a new SDEV_TRANSPORT_OFFLINE state has been added to help distinguish transport problems from another I/O failure causes. Transport errors are now handled differently and storage devices can now recover from these failures without user intervention.

BZ#880085

Previously, the IP over Infiniband (IPoIB) driver maintained state information about neighbors on the network by attaching it to the core network's neighbor structure. However, due to a race condition between the freeing of the core network neighbor struct and the freeing of the IPoIB network struct, a use after free condition could happen, resulting in either a kernel oops or 4 or 8 bytes of kernel memory being zeroed when it was not supposed to be. These patches decouple the IPoIB neighbor struct from the core networking stack's neighbor struct so that there is no race between the freeing of one and the freeing of the other.

BZ#880928

When a new rpc_task is created, the code takes a reference to rpc_cred and sets the task->tk_cred pointer to it. After the call completes, the resources held by the rpc_task are freed. Previously, however, after the rpc_cred was released, the pointer to it was not zeroed out. This led to an rpc_cred reference count underflow, and consequently to a kernel panic. With this update, the pointer to rpc_cred is correctly zeroed out, which prevents a kernel panic from occurring in this scenario.

BZ#884422

Previously, the HP Smart Array driver (hpsa) used the target reset functionality. However, HP Smart Array logical drives do not support the target reset functionality. Therefore, if the target reset failed, the logical drive was taken offline with a file system error. The hpsa driver has been updated to use the LUN reset functionality instead of target reset, which is supported by these drives.

BZ#886618

The bonding driver previously did not honor the maximum Generic Segmentation Offload (GSO) length of packets and segments requested by the underlying network interface. This caused the firmware of the underlying NIC, such as be2net, to become unresponsive. This update modifies the bonding driver to set up the lowest gso_max_size and gso_max_segs values of network devices while attaching and detaching the devices as slaves. The network drivers no longer hangs and network traffic now proceeds as expected in setups using a bonding interface.

BZ#886760

Previously, the interrupt handlers of the qla2xxx driver could clear pending interrupts right after the IRQ lines were attached during system start-up. Consequently, the kernel could miss the interrupt that reported completion of the link initialization, and the qla2xxx driver then failed to detect all attached LUNs. With this update, the qla2xxx driver has been modified to no longer clear interrupt bits after attaching the IRQ lines. The driver now correctly detects all attached LUNs as expected.

BZ#888215

When TCP segment offloading (TSO) or jumbo packets are used on the Broadcom BCM5719 network interface controller (NIC) with multiple TX rings, small packets can be starved for resources by the simple round-robin hardware scheduling of these TX rings, thus causing lower network performance. To ensure reasonable network performance for all NICs, multiple TX rings are now disabled by default.

BZ#888818

Due to insufficient handling of a dead Input/Output Controller (IOC), the mpt2sas driver could fail Enhanced I/O Error Handling (EEH) recovery for certain PCI bus failures on 64-bit IBM PowerPC machines. With this update, when a dead IOC is detected, EEH recovery routine has more time to resolve the failure and the controller in a non-operational state is removed.

BZ#891580

A possible race between the n_tty_read() and reset_buffer_flags() functions could result in a NULL pointer dereference in the n_tty_read() function under certain circumstances. As a consequence, a kernel panic could have been triggered when interrupting a current task on a serial console. This update modifies the tty driver to use a spin lock to prevent functions from a parallel access to variables. A NULL pointer dereference causing a kernel panic can no longer occur in this scenario.

All users should upgrade to these updated packages, which contain backported patches to correct these issues and bugs. The system must be rebooted for this update to take effect.

5.135.9. RHSA-2012:1064 – Important: kernel security and bug fix update

Updated kernel packages that fix two 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 score, which gives a detailed severity rating, is 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-2012-2744, Important

A NULL pointer dereference flaw was found in the nf_ct_frag6_reasm() function in the Linux kernel's netfilter IPv6 connection tracking implementation. A remote attacker could use this flaw to send specially-crafted packets to a target system that is using IPv6 and also has the nf_conntrack_ipv6 kernel module loaded, causing it to crash.

CVE-2012-2745, Moderate

A flaw was found in the way the Linux kernel's key management facility handled replacement session keyrings on process forks. A local, unprivileged user could use this flaw to cause a denial of service.

Red Hat would like to thank an anonymous contributor working with the Beyond Security SecuriTeam Secure Disclosure program for reporting CVE-2012-2744.

Bug Fixes

BZ#832359

Previously introduced firmware files required for new Realtek chipsets contained an invalid prefix ("rtl_nic_") in the file names, for example "/lib/firmware/rtl_nic/rtl_nic_rtl8168d-1.fw". This update corrects these file names. For example, the aforementioned file is now correctly named "/lib/firmware/rtl_nic/rtl8168d-1.fw".

BZ#832363

This update blacklists the ADMA428M revision of the 2GB ATA Flash Disk device. This is due to data corruption occurring on the said device when the Ultra-DMA 66 transfer mode is used. When the "libata.force=5:pio0,6:pio0" kernel parameter is set, the aforementioned device works as expected.

BZ#832365

On Red Hat Enterprise Linux 6, mounting an NFS export from a Windows 2012 server failed due to the fact that the Windows server contains support for the minor version 1 (v4.1) of the NFS version 4 protocol only, along with support for versions 2 and 3. The lack of the minor version 0 (v4.0) support caused Red Hat Enterprise Linux 6 clients to fail instead of rolling back to version 3 as expected. This update fixes this bug and mounting an NFS export works as expected.

BZ#833034

On ext4 file systems, when fallocate() failed to allocate blocks due to the ENOSPC condition (no space left on device) for a file larger than 4 GB, the size of the file became corrupted and, consequently, caused file system corruption. This was due to a missing cast operator in the "ext4_fallocate()" function. With this update, the underlying source code has been modified to address this issue, and file system corruption no longer occurs.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

5.135.10. RHBA-2012:1199 – kernel bug fix update

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

The kernel packages contain the Linux kernel, the core of any Linux operating system.

When an NTP server asserts the STA_INS flag (Leap Second Insert), the kernel starts an hrtimer (high-resolution timer) with a countdown clock. This hrtimer expires at end of the current month, midnight UTC, and inserts a second into the kernel timekeeping structures. A scheduled leap second occurred on June 30 2012 midnight UTC.

Bug Fixes

BZ#840950

Previously in the kernel, when the leap second hrtimer was started, it was possible that the kernel livelocked on the xtime_lock variable. This update fixes the problem by using a mixture of separate subsystem locks (timekeeping and ntp) and removing the xtime_lock variable, thus avoiding the livelock scenarios that could occur in the kernel.

BZ#847366

After the leap second was inserted, applications calling system calls that used futexes consumed almost 100% of available CPU time. This occurred because the kernel's timekeeping structure update did not properly update these futexes. The futexes repeatedly expired, re-armed, and then expired

immediately again. This update fixes the problem by properly updating the futex expiration times by calling the clock_was_set_delayed() function, an interrupt-safe method of the clock_was_set() function.

All users are advised to upgrade to these updated packages, which fix these bugs. The system must be rebooted for this update to take effect.

5.135.11. RHSA-2012:1156 – Moderate: kernel security and bug fix update

Updated kernel packages that fix two 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 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 kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2012-2383, Moderate

An integer overflow flaw was found in the i915_gem_execbuffer2() function in the Intel i915 driver in the Linux kernel. A local, unprivileged user could use this flaw to cause a denial of service. This issue only affected 32-bit systems.

CVE-2011-1078, Low

A missing initialization flaw was found in the sco_sock_getsockopt_old() function in the Linux kernel's Bluetooth implementation. A local, unprivileged user could use this flaw to cause an information leak.

Red Hat would like to thank Vasiliy Kulikov of Openwall for reporting the CVE-2011-1078 issue.

Bug Fixes

BZ#832360

A bug in the writeback livelock avoidance scheme could result in some dirty data not being written to disk during a sync operation. In particular, this could occasionally occur at unmount time, when previously written file data was not synced, and was unavailable after the file system was remounted. Patches have been applied to address this issue, and all dirty file data is now synced to disk at unmount time.

BZ#838821

During the update of the be2net driver between the Red Hat Enterprise Linux 6.1 and 6.2, the NETIF_F_GRO flag was incorrectly removed, and the GRO (Generic Receive Offload) feature was therefore disabled by default. In OpenVZ kernels based on Red Hat Enterprise Linux 6.2, this led to a significant traffic decrease. To prevent this problem, the NETIF_F_GRO flag has been included in the underlying source code.

BZ#840023

Previously, the size of the multicast IGMP (Internet Group Management Protocol) snooping hash table for a bridge was limited to 256 entries even though the maximum is 512. This was due to the hash table size being incorrectly compared to the maximum hash table size, hash_max, and the

following message could have been produced by the kernel:

> Multicast hash table maximum reached, disabling snooping: vnet1, 512

With this update, the hash table value is correctly compared to the hash_max value, and the error message no longer occurs under these circumstances.

BZ#840052

In the ext4 file system, splitting an unwritten extent while using Direct I/O could fail to mark the modified extent as dirty, resulting in multiple extents claiming to map the same block. This could lead to the kernel or fsck reporting errors due to multiply claimed blocks being detected in certain inodes. In the ext4_split_unwritten_extents() function used for Direct I/O, the buffer which contains the modified extent is now properly marked as dirty in all cases. Errors due to multiply claimed blocks in inodes should no longer occur for applications using Direct I/O.

BZ#840156

With certain switch peers and firmware, excessive link flaps could occur due to the way DCBX (Data Center Bridging Exchange) was handled. To prevent link flaps, changes were made to examine the capabilities in more detail and only initialize hardware if the capabilities have changed.

BZ#841406

The CONFIG_CFG80211_WEXT configuration option previously defined in the KConfig of the ipw2200 driver was removed with a recent update. This led to a build failure of the driver. The driver no longer depends on the CONFIG_CFG80211_WEXT option, so it can build successfully.

BZ#841411

Migrating virtual machines from Intel hosts that supported the VMX "Unrestricted Guest" feature to older hosts without this feature could result in kvm returning the "unhandled exit 80000021" error for guests in real mode. The underlying source code has been modified so that migration completes successfully on hosts where "Unrestricted Guest" is disabled or not supported.

BZ#841579

Previous update changed the /proc/stat code to use the get_cpu_idle_time_us() and get_cpu_iowait_time_us() macros if dynamic ticks are enabled in the kernel. This could lead to problems on IBM System z architecture that defines the arch_idle_time() macro. For example, executing the "vmstat" command could fail with "Floating point exception" followed by a core dump. The underlying source code has been modified so that the arch_idle_time() macro is used for idle and iowait times, which prevents the mentioned problem.

BZ#842429

Bond masters and slaves now have separate VLAN groups. As such, if a slave device incurred a network event that resulted in a failover, the VLAN device could process this event erroneously. With this update, when a VLAN is attached to a master device, it ignores events generated by slave devicec so that the VLANs do not go down until the bond master does.

Users should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

5.135.12. RHSA-2012:1580 – Moderate: kernel security, bug fix and enhancement update

Updated kernel packages that fix multiple security issues, numerous 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.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2012-2375, Moderate

It was found that the RHSA-2012:0862 update did not correctly fix the CVE-2011-4131 issue. A malicious Network File System version 4 (NFSv4) server could return a crafted reply to a GETACL request, causing a denial of service on the client.

CVE-2012-4565, Moderate

A divide-by-zero flaw was found in the TCP Illinois congestion control algorithm implementation in the Linux kernel. If the TCP Illinois congestion control algorithm were in use (the sysctl net.ipv4.tcp_congestion_control variable set to "illinois"), a local, unprivileged user could trigger this flaw and cause a denial of service.

CVE-2012-5517, Moderate

A NULL pointer dereference flaw was found in the way a new node's hot added memory was propagated to other nodes' zonelists. By utilizing this newly added memory from one of the remaining nodes, a local, unprivileged user could use this flaw to cause a denial of service.

CVE-2012-2100, Low

It was found that the initial release of Red Hat Enterprise Linux 6 did not correctly fix the CVE-2009-4307 issue, a divide-by-zero flaw in the ext4 file system code. A local, unprivileged user with the ability to mount an ext4 file system could use this flaw to cause a denial of service.

CVE-2012-4444, Low

A flaw was found in the way the Linux kernel's IPv6 implementation handled overlapping, fragmented IPv6 packets. A remote attacker could potentially use this flaw to bypass protection mechanisms (such as a firewall or intrusion detection system (IDS)) when sending network packets to a target system.

Red Hat would like to thank Antonios Atlasis working with Beyond Security's SecuriTeam Secure Disclosure program and Loganaden Velvindron of AFRINIC for reporting CVE-2012-4444. The CVE-2012-2375 issue was discovered by Jian Li of Red Hat, and CVE-2012-4565 was discovered by Rodrigo Freire of Red Hat.

Bug Fixes

BZ#853950

The kernel allows high priority real time tasks, such as tasks scheduled with the SCHED_FIFO policy, to be throttled. Previously, the CPU stop tasks were scheduled as high priority real time tasks and could be thus throttled accordingly. However, the replenishment timer, which is responsible for clearing a throttle flag on tasks, could be pending on the just disabled CPU. This could lead to a situation that the throttled tasks were never scheduled to run. Consequently, if any of such tasks was needed to complete the CPU disabling, the system became unresponsive. This update introduces a

new scheduler class, which gives a task the highest possible system priority and such a task cannot be throttled. The stop-task scheduling class is now used for the CPU stop tasks, and the system shutdown completes as expected in the scenario described.

BZ#864826

A kernel panic occurred when the size of a block device was changed and an I/O operation was issued at the same time. This was because the direct and non-direct I/O code was written with the assumption that the block size would not change. This update introduces a new read-write lock, bd_block_size_semaphore. The lock is taken for read during I/O operations and for write when changing the block size of a device. As a result, block size cannot be changed while I/O is being submitted. This prevents the kernel from crashing in the described scenario.

BZ#866470

A previous kernel update introduced a bug that caused RAID0 and linear arrays larger than 4 TB to be truncated to 4 TB when using 0.90 metadata. The underlying source code has been modified so that 0.90 RAID0 and linear arrays larger than 4 TB are no longer truncated in the md RAID layer.

BZ#866795

The mlx4 driver must program the mlx4 card so that it is able to resolve which MAC addresses to listen to, including multicast addresses. Therefore, the mlx4 card keeps a list of trusted MAC addresses. The driver used to perform updates to this list on the card by emptying the entire list and then programming in all of the addresses. Thus, whenever a user added or removed a multicast address or put the card into or out of promiscuous mode, the card's entire address list was re-written. This introduced a race condition, which resulted in a packet loss if a packet came in on an address the card should be listening to, but had not yet been reprogrammed to listen to. With this update, the driver no longer rewrites the entire list of trusted MAC addresses on the card but maintains a list of addresses that are currently programmed into the card. On address addition, only the new address is added to the end of the list, and on removal, only the to-be-removed address is removed from the list. The mlx4 card no longer experiences the described race condition and packets are no longer dropped in this scenario.

BZ#871854

If there are no active threads using a semaphore, blocked threads should be unblocked. Previously, the R/W semaphore code looked for a semaphore counter as a whole to reach zero - which is incorrect because at least one thread is usually queued on the semaphore and the counter is marked to reflect this. As a consequence, the system could become unresponsive when an application used direct I/O on the XFS file system. With this update, only the count of active semaphores is checked, thus preventing the hang in this scenario.

BZ#874022

Due to an off-by-one error in a test condition in the bnx2x_start_xmit and bnx2x_tx_int functions, the TX queue of a NIC could, under some circumstances, be prevented from being resumed. Consequently, NICs using the bnx2x driver, such as Broadcom NetXtreme II 10G network devices, went offline. To bring the NIC back online, the bnx2x module had to be reloaded. This update corrects the test condition in the mentioned functions and the NICs using the bnx2x driver work as expected in the described scenario.

BZ#876088

If an abort request times out to the virtual Fibre Channel adapter, the ibmvfc driver initiates a reset of the adapter. Previously, however, the ibmvfc driver incorrectly returned success to the eh_abort handler and then sent a response to the same command, which led to a kernel oops on IBM System p

machines. This update ensures that both the abort request and the request being aborted are completed prior to exiting the en_abort handler, and the kernel oops no longer occurs in this scenario.

BZ#876101

The hugetlbfs file system implementation was missing a proper lock protection of enqueued huge pages at the gather_surplus_pages() function. Consequently, the hstate.hugepages_freelist list became corrupted, which caused a kernel panic. This update adjusts the code so that the used spinlock protection now assures atomicity and safety of enqueued huge pages when handling hstate.hugepages_freelist. The kernel no longer panics in this scenario.

BZ#876487

A larger command descriptor block (CDB) is allocated for devices using Data Integrity Field (DIF) type 2 protection. The CDB was being freed in the sd_done() function, which resulted in a kernel panic if the command had to be retried in certain error recovery cases. With this update, the larger CDB is now freed in the sd_unprep_fn() function instead. This prevents the kernel panic from occurring.

BZ#876491

The previous implementation of socket buffers (SKBs) allocation for a NIC was node-aware, that is, memory was allocated on the node closest to the NIC. This increased performance of the system because all DMA transfer was handled locally. This was a good solution for networks with a lower frame transmitting rate where CPUs of the local node handled all the traffic of the single NIC well. However, when using 10Gb Ethernet devices, CPUs of one node usually do not handle all the traffic of a single NIC efficiently enough. Therefore, system performance was poor even though the DMA transfer was handled by the node local to the NIC. This update modifies the kernel to allow SKBs to be allocated on a node that runs applications receiving the traffic. This ensures that the NIC's traffic is handled by as many CPUs as needed, and since SKBs are accessed very frequently after allocation, the kernel can now operate much more efficiently even though the DMA can be transferred cross-node.

BZ#876493

When performing PCI device assignment on AMD systems, a virtual machine using the assigned device could not be able to boot, as the device had failed the assignment, leaving the device in an unusable state. This was due to an improper range check that omitted the last PCI device in a PCI subsystem or tree. The check has been fixed to include the full range of PCI devices in a PCI subsystem or tree. This bug fix avoids boot failures of a virtual machine when the last device in a PCI subsystem is assigned to a virtual machine on an AMD host system.

BZ#876496

The mmap_rnd() function is expected to return a value in the [0x00000000 .. 0x000FF000] range on 32-bit x86 systems. This behavior is used to randomize the base load address of shared libraries by a bug fix resolving the CVE-2012-1568 issue. However, due to a signedness bug, the mmap_rnd() function could return values outside of the intended scope. Consequently, the shared libraries base address could be less than one megabyte. This could cause binaries that use the MAP_FIXED mappings in the first megabyte of the process address space (typically, programs using vm86 functionality) to work incorrectly. This update modifies the mmap_rnd() function to no longer cast values returned by the get_random_int() function to the long data type. The aforementioned binaries now work correctly in this scenario.

BZ#876499

Previously, XFS could, under certain circumstances, incorrectly read metadata from the journal

during XFS log recovery. As a consequence, XFS log recovery terminated with an error message and prevented the file system from being mounted. This problem could result in a loss of data if the user forcibly "zeroed" the log to allow the file system to be mounted. This update ensures that metadata is read correctly from the log so that journal recovery completes successfully and the file system mounts as expected.

BZ#876549

Some BIOS firmware versions could leave the "Frame Start Delay" bits of the PIPECONF register in test mode on selected Intel chipsets. Consequently, video output on certain Lenovo laptop series, such as T41x or T42x, became corrupted (for example, the screen appeared to be split and shifted to the right) after upgrading VBIOS from version 2130 to 2132. This update corrects the problem by resetting the "Frame Start Delay" bits for the normal operation use in the DRM driver. Video output of the previously affected Lenovo models is now correct.

Enhancement

BZ#877950

The INET socket interface has been modified to send a warning message when the ip_options structure is allocated directly by a third-party module using the kmalloc() function.

Users should upgrade to these updated kernel packages, which contain backported patches to correct these issues, fix these bugs and add this enhancement. The system must be rebooted for this update to take effect.

5.135.13. RHSA-2012:1426 – Moderate: 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 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 kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2012-2133, Moderate

A use-after-free flaw was found in the Linux kernel's memory management subsystem in the way quota handling for huge pages was performed. A local, unprivileged user could use this flaw to cause a denial of service or, potentially, escalate their privileges.

CVE-2012-3511, Moderate

A use-after-free flaw was found in the madvise() system call implementation in the Linux kernel. A local, unprivileged user could use this flaw to cause a denial of service or, potentially, escalate their privileges.

CVE-2012-1568, Low

It was found that when running a 32-bit binary that uses a large number of shared libraries, one of the libraries would always be loaded at a predictable address in memory. An attacker could use this flaw to bypass the Address Space Layout Randomization (ASLR) security feature.

CVE-2012-3400, **Low**

Buffer overflow flaws were found in the udf_load_logicalvol() function in the Universal Disk Format (UDF) file system implementation in the Linux kernel. An attacker with physical access to a system could use these flaws to cause a denial of service or escalate their privileges.

Red Hat would like to thank Shachar Raindel for reporting CVE-2012-2133.

Bug Fixes

BZ#865713

Previously, the I/O watchdog feature was disabled when Intel Enhanced Host Controller Interface (EHCI) devices were detected. This could cause incorrect detection of USB devices upon addition or removal. Also, in some cases, even though such devices were detected properly, they were non-functional. The I/O watchdog feature can now be enabled on the kernel command line, which improves hardware detection on underlying systems.

BZ#864821

The usb_device_read() routine used the bus->root_hub pointer to determine whether or not the root hub was registered. However, this test was invalid because the pointer was set before the root hub was registered and remained set even after the root hub was unregistered and deallocated. As a result, the usb_device_read() routine accessed freed memory, causing a kernel panic; for example, on USB device removal. With this update, the hcs->rh_registered flag - which is set and cleared at the appropriate times - is used in the test, and the kernel panic no longer occurs in this scenario.

BZ#853257

Previously, when a server attempted to shut down a socket, the svc_tcp_sendto() function set the XPT_CLOSE variable if the entire reply failed to be transmitted. However, before XPT_CLOSE could be acted upon, other threads could send further replies before the socket was really shut down. Consequently, data corruption could occur in the RPC record marker. With this update, send operations on a closed socket are stopped immediately, thus preventing this bug.

BZ#853943

Previously, a race condition existed whereby device open could race with device removal (for example when hot-removing a storage device), potentially leading to a kernel panic. This was due a use-after-free error in the block device open patch, which has been corrected by not referencing the "disk" pointer after it has been passed to the module_put() function.

BZ#854476

Sometimes, the crypto allocation code could become unresponsive for 60 seconds or multiples thereof due to an incorrect notification mechanism. This could cause applications, like openswan, to become unresponsive. The notification mechanism has been improved to avoid such hangs.

BZ#856106

Traffic to the NFS server could trigger a kernel oops in the svc_tcp_clear_pages() function. The source code has been modified, and the kernel oops no longer occurs in this scenario.

BZ#860784

When a device was registered to a bus, a race condition could occur between the device being added to the list of devices of the bus and binding the device to a driver. As a result, the device could already be bound to a driver which led to a warning and incorrect reference counting, and

consequently to a kernel panic on device removal. To avoid the race condition, this update adds a check to identify an already bound device.

BZ#865308

When I/O is issued through blk_execute_rq(), the blk_execute_rq_nowait() routine is called to perform various tasks. At first, the routine checks for a dead queue. Previously, however, if a dead queue was detected, the blk_execute_rq_nowait() function did not invoke the done() callback function. This resulted in blk_execute_rq() being unresponsive when waiting for completion, which had never been issued. To avoid such hangs, the rq->end_io pointer is initialized to the done() callback before the queue state is verified.

BZ#860942

The Out of Memory (OOM) killer killed processes outside of a memory cgroup when one or more processes inside that memory cgroup exceeded the "memory.limit_in_bytes" value. This was because when a copy-on-write fault happened on a Transparent Huge Page (THP), the 2 MB THP caused the cgroup to exceed the memory.limit_in_bytes value but the individual 4 KB page was not exceeded. With this update, the 2 MB THP is correctly split into 4 KB pages when the memory.limit_in_bytes value is exceeded. The OOM kill is delivered within the memory cgroup; tasks outside the memory cgroups are no longer killed by the OOM killer.

BZ#857055

An unnecessary check for the RXCW.CW bit could cause the Intel e1000e NIC (Network Interface Controller) to not work properly. The check has been removed so that the Intel e1000e NIC now works as expected.

BZ#860640

A kernel oops could occur due to a NULL pointer dereference upon USB device removal. The NULL pointer dereference has been fixed and the kernel no longer crashes in this scenario.

BZ#864827

Previously, a use-after-free bug in the usbhid code caused a NULL pointer dereference. Consequent kernel memory corruption resulted in a kernel panic and could cause data loss. This update adds a NULL check to avoid these problems.

BZ#841667

USB Request Blocks (URBs) coming from user space were not allowed to have transfer buffers larger than an arbitrary maximum. This could lead to various problems; for example, attempting to redirect certain USB mass-storage devices could fail. To avoid such problems, programs are now allowed to submit URBs of any size; if there is not sufficient contiguous memory available, the submission fails with an ENOMEM error. In addition, to prevent programs from submitting a lot of small URBs and so using all the DMA-able kernel memory, this update also replaces the old limits on individual transfer buffers with a single global limit of 16MB on the total amount of memory in use by USB file system (usbfs).

BZ#841824

A USB Human Interface Device (HID) can be disconnected at any time. If this happened right before or while the hiddev_ioctl() call was in progress, hiddev_ioctl() attempted to access the invalid hiddev->hid pointer. When the HID device was disconnected, the hiddev_disconnect() function called the hid_device_release() function, which frees the hid_device structure type, but did not set the hiddev->hid pointer to NULL. If the deallocated memory region was re-used by the kernel, a kernel panic or

memory corruption could occur. The hiddev->exist flag is now checked while holding the existancelock and hid_device is used only if such a device exists. As a result, the kernel no longer crashes in this scenario.

BZ#863147

The MAC address stored in the driver's private structure is of the unsigned character data type but parameters of the strlcpy() function are of the signed character data type. This conversion of data types led to change in the value. This changed value was passed to the upper layer and garbage characters were displayed when running the "iscsiadm -m iface" command. Consequently, the garbage characters in the MAC address led to boot failures of iSCSI devices. MAC addresses are now formatted using the sysfs_format_mac() function rather than strlcpy(), which prevents the described problems.

BZ#861953

It is possible to receive data on multiple transports. Previously, however, data could be selectively acknowledged (SACKed) on a transport that had never received any data. This was against the SHOULD requirement in section 6.4 of the RFC 2960 standard. To comply with this standard, bundling of SACK operations is restricted to only those transports which have moved the ctsn of the association forward since the last sack. As a result, only outbound SACKs on a transport that has received a chunk since the last SACK are bundled.

BZ#861390

Bugs in the lpfs driver caused disruptive logical unit resets during fabric fault testing. The underlying source code has been modified so that the problem no longer occurs.

BZ#852450

Previously, bnx2x devices did not disable links with a large number of RX errors and overruns, and such links could still be detected as active. This prevented the bonding driver from failing over to a working link. This update restores remote-fault detection, which periodically checks for remote faults on the MAC layer. In case the physical link appears to be up but an error occurs, the link is disabled. Once the error is cleared, the link is brought up again.

BZ#860787

Various race conditions that led to indefinite log reservation hangs due to xfsaild "idle" mode occurred in XFS file system. This could lead to certain tasks being unresponsive; for example, the cp utility could become unresponsive on heavy workload. This update improves the Active Item List (AIL) pushing logic in xfsaild. Also, the log reservation algorithm and interactions with xfsaild have been improved. As a result, the aforementioned problems no longer occur in this scenario.

BZ#858955

On dual port Mellanox hardware, the mlx4 driver was adding promiscuous mode to the correct port, but when attempting to remove promiscuous mode from a port, it always tried to remove it from port one. It was therefore impossible to remove promiscuous mode from the second port, and promiscuous mode was incorrectly removed from port one even if it was not intended. With this update, the driver now properly attempts to remove promiscuous mode from port two when needed.

BZ#858956

Mellanox hardware keeps a separate list of Ethernet hardware addresses it listens to depending on whether the Ethernet hardware address is unicast or multicast. Previously, the mlx4 driver was incorrectly adding multicast addresses to the unicast list. This caused unstable behavior in terms of whether or not the hardware would have actually listened to the addresses requested. This update fixes the problem by always putting multicast addresses on the multicast list and vice versa.

BZ#859326

If a dirty GFS2 inode was being deleted but was in use by another node, its metadata was not written out before GFS2 checked for dirty buffers in the gfs2_ail_flush() function. GFS2 was relying on the inode_go_sync() function to write out the metadata when the other node tried to free the file. However, this never happened because GFS2 failed the error check. With this update, the inode is written out before calling the gfs2_ail_flush() function. If a process has the PF_MEMALLOC flag set, it does not start a new transaction to update the access time when it writes out the inode. The inode is marked as dirty to make sure that the access time is updated later unless the inode is being freed.

BZ#859436

In a previous release of Red Hat Enterprise Linux, the new Mellanox packet steering architecture had been intentionally left out of the Red Hat kernel. With Red Hat Enterprise Linux 6.3, the new Mellanox packet steering architecture was merged into Red Hat Mellanox driver. One merge detail was missing, and as a result, the multicast promiscuous flag on an interface was not checked during an interface reset to see if the flag was on prior to the reset and should be re-enabled after the reset. This update fixes the problem, so if an adapter is reset and the multicast promiscuous flag was set prior to the reset, the flag is now still set after the reset.

BZ#860165

Previously, the default minimum entitled capacity of a virtual processor was 10%. This update changes the PowerPC architecture vector to support a lower minimum virtual processor capacity of 1%.

BZ#858954

Previously, a cgroup or its hierarchy could only be modified under the cgroup_mutex master lock. This introduced a locking dependency on cred_guard_mutex from cgroup_mutex and completed a circular dependency, which involved cgroup_mutex, namespace_sem and workqueue, and led to a deadlock. As a consequence, many processes were unresponsive, and the system could be eventually unusable. This update introduces a new mutex, cgroup_root_mutex, which protects cgroup root modifications and is now used by mount options readers instead of the master lock. This breaks the circular dependency and avoids the deadlock.

All users of kernel should upgrade to these updated packages, which contain backported patches to correct these issues. The system must be rebooted for this update to take effect.

5.135.14. RHSA-2012:0862 – Moderate: Red Hat Enterprise Linux 6.3 kernel security, bug fix, and enhancement update

Updated kernel packages that fix two 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 third regular update.

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 kernel packages contain the Linux kernel, the core of any Linux operating system.

Security Fixes

CVE-2011-1083, **Moderate**

A flaw was found in the way the Linux kernel's Event Poll (epoll) subsystem handled large, nested epoll structures. A local, unprivileged user could use this flaw to cause a denial of service.

CVE-2011-4131, **Moderate**

A malicious Network File System version 4 (NFSv4) server could return a crafted reply to a GETACL request, causing a denial of service on the client.

Red Hat would like to thank Nelson Elhage for reporting CVE-2011-1083, and Andy Adamson for reporting CVE-2011-4131.

Bug Fixes

BZ#824025

Hotplugging SATA disks did not work properly and the system experienced various issues when hotplugging such devices. This update fixes several hotplugging issues in the kernel and SAS hotplugging now works as expected.

BZ#782374

Due to a bug in the hid_reset() function, a deadlock could occur when a Dell iDRAC controller was reset. Consequently, its USB keyboard or mouse device became unresponsive. A patch that fixes the underlying code has been provided to address this bug and the hangs no longer occur in the described scenario.

BZ#781531

The AMD IOMMU driver used wrong shift direction in the alloc_new_range() function. Consequently, the system could terminate unexpectedly or become unresponsive. This update fixes the code and crashes and hangs no longer occur in the described scenario.

BZ#781524

Previously, the AMD IOMMU (input/output memory management unit) driver could use the MSI address range for DMA (direct memory access) addresses. As a consequence, DMA could fail and spurious interrupts would occur if this address range was used. With this update, the MSI address range is reserved to prevent the driver from allocating wrong addresses and DMA is now assured to work as expected in the described scenario.

BZ#773705

Windows clients never send write requests larger than 64 KB but the default size for write requests in Common Internet File System (CIFS) was set to a much larger value. Consequently, write requests larger than 64 KB caused various problems on certain third-party servers. This update lowers the default size for write requests to prevent this bug. The user can override this value to a larger one to get better performance.

BZ#773522

Due to a race condition between the notify_on_release() function and task movement between cpuset or memory cgroup directories, a system deadlock could occur. With this update, the cgroup_wq cgroup has been created and both async_rebuild_domains() and check_for_release() functions used for task movements use it, thus fixing this bug.

BZ#773517

Due to invalid calculations of the vruntime variable along with task movement between cgroups, moving tasks between cgroups could cause very long scheduling delays. This update fixes this problem by setting the cfs_rq and curr parameters after holding the rq->lock lock.

BZ#784671

The kernel code checks for conflicts when an application requests a specific port. If there is no conflict, the request is granted. However, the port auto-selection done by the kernel failed when all ports were bound, even if there was an available port with no conflicts. With this update, the port auto-selection code has been fixed to properly use ports with no conflicts.

BZ#784758

A bug in the try_to_wake_up() function could cause status change from TASK_DEAD to TASK_RUNNING in a race condition with an SMI (system management interrupt) or a guest environment of a virtual machine. As a consequence, the exited task was scheduled again and a kernel panic occurred. This update fixes the race condition in the do_exit() function and the panic no longer occurs in the described scenario.

BZ#785891

Previously, if more than a certain number of qdiscs (Classless Queuing Disciplines) using the autohandle mechanism were allocated a soft lock-up error occurred. This update fixes the maximum loop count and adds the cond_resched() call in the loop, thus fixing this bug.

BZ#785959

Prior to this update, the find_busiest_group() function used sched_group->cpu_power in the denominator of a fraction with a value of 0. Consequently, a kernel panic occurred. This update prevents the divide by zero in the kernel and the panic no longer occurs.

BZ#772874

In the Common Internet File System (CIFS), the oplock break jobs and async callback handlers both use the SLOW-WORK workqueue, which has a finite pool of threads. Previously, these oplock break jobs could end up taking all the running queues waiting for a page lock which blocks the callback required to free this page lock from being completed. This update separates the oplock break jobs into a separate workqueue VERY-SLOW-WORK, allowing the callbacks to be completed successfully and preventing the deadlock.

BZ#772317

Previously, network drivers that had Large Receive Offload (LRO) enabled by default caused the system to run slow, lose frame, and eventually prevent communication, when using software bridging. With this update, LRO is automatically disabled by the kernel on systems with a bridged configuration, thus preventing this bug.

BZ#772237

When transmitting a fragmented socket buffer (SKB), the qlge driver fills a descriptor with fragment addresses, after DMA-mapping them. On systems with pages larger than 8 KB and less than eight fragments per SKB, a macro defined the size of the OAL (Outbound Address List) list as 0. For SKBs with more than eight fragments, this would start overwriting the list of addresses already mapped and would make the driver fail to properly unmap the right addresses on architectures with pages larger than 8 KB. With this update, the size of external list for TX address descriptors have been fixed and qlge no longer fails in the described scenario.

BZ#772136

Prior to this update, the wrong size was being calculated for the vfinfo structure. Consequently, networking drivers that created a large number of virtual functions caused warning messages to appear when loading and unloading modules. Backported patches from upstream have been provided to resolve this issue, thus fixing this bug.

BZ#771251

The fcoe_transport_destroy path uses a work queue to destroy the specified FCoE interface. Previously, the destroy_work work queue item blocked another single-threaded work queue. Consequently, a deadlock between queues occurred and the system became unresponsive. With this update, fcoe_transport_destroy has been modified and is now a synchronous operation, allowing to break the deadlock dependency. As a result, destroy operations are now able to complete properly, thus fixing this bug.

BZ#786518

On a system that created and deleted lots of dynamic devices, the 31-bit Linux ifindex object failed to fit in the 16-bit macvtap minor range, resulting in unusable macvtap devices. The problem primarily occurred in a libvirt-controlled environment when many virtual machines were started or restarted, and caused libvirt to report the following message: Error starting domain: cannot open macvtap tap device /dev/tap222364: No such device or address With this update, the macvtap's minor device number allocation has been modified so that virtual machines can now be started and restarted as expected in the described scenario.

BZ#770023

A bug in the splice code has caused the file position on the write side of the sendfile() system call to be incorrectly set to the read side file position. This could result in the data being written to an incorrect offset. Now, sendfile() has been modified to correctly use the current file position for the write side file descriptor, thus fixing this bug.

Note that in the following common sendfile() scenarios, this bug does not occur: when both read and write file positions are identical and when the file position is not important (e.g. if the write side is a socket).

BZ#769626

Prior to this update, Active State Power Management (ASPM) was not properly disabled, and this interfered with the correct operation of the hpsa driver. Certain HP BIOS versions do not report a proper disable bit, and when the kernel fails to read this bit, the kernel defaults to enabling ASPM. Consequently, certain servers equipped with a HP Smart Array controller were unable to boot unless the pcie_aspm=off option was specified on the kernel command line. A backported patch has been provided to address this problem, ASPM is now properly disabled, and the system now boots up properly in the described scenario.

BZ#769590

Due to a race condition, running the "ifenslave -d bond0 eth0" command to remove the slave interface from the bonding device could cause the system to crash when a networking packet was being received at the same time. With this update, the race condition has been fixed and the system no longer crashes under these circumstances.

BZ#769007

In certain circumstances, the qla2xxx driver was unable to discover fibre channel (FC) tape devices because the ADISC ELS request failed. This update adds the new module parameter, ql2xasynclogin, to address this issue. When this parameter is set to "0", FC tape devices are discovered properly.

BZ#786960

When running AF_IUCV socket programs with IUCV transport, an IUCV SEVER call was missing in the callback of a receiving IUCV SEVER interrupt. Under certain circumstances, this could prevent z/VM from removing the corresponding IUCV-path completely. This update adds the IUCV SEVER call to the callback, thus fixing this bug. In addition, internal socket states have been merged, thus simplifying the AF_IUCV code.

BZ#767753

When the nohz=off kernel parameter was set, kernel could not enter any CPU C-state. With this update, the underlying code has been fixed and transitions to CPU idle states now work as expected.

BZ#766861

Under heavy memory and file system load, the "mapping->nrpages == 0" assertion could occur in the end_writeback() function. As a consequence, a kernel panic could occur. This update provides a reliable check for mapping->nrpages that prevent the described assertion, thus fixing this bug.

BZ#765720

An insufficiently designed calculation in the CPU accelerator in previous kernel caused an arithmetic overflow in the sched_clock() function when system uptime exceeded 208.5 days. This overflow led to a kernel panic on the systems using the Time Stamp Counter (TSC) or Virtual Machine Interface (VMI) clock source. This update corrects the aforementioned calculation so that this arithmetic overflow and kernel panic can no longer occur under these circumstances.

BZ#765673

Previously, the cfq_cic_link() function had a race condition. When some processes, which shared ioc issue I/O to the same block device simultaneously, cfq_cic_link() sometimes returned the -EEXIST error code. Consequently, one of the processes started to wait indefinitely. A patch has been provided to address this issue and the cfq_cic_lookup() call is now retried in the described scenario, thus fixing this bug.

BZ#667925

Previously, the SFQ qdisc packet scheduler class had no bind_tcf() method. Consequently, if a filter was added with the classid parameter to SFQ, a kernel panic occurred due to a null pointer dereference. With this update, the dummy ".unbind_tcf" and ".put" qdisc class options have been added to conform with the behaviour of other schedulers, thus fixing this bug.

BZ#787762

Previously, an incorrect portion of memory was freed when unmapping a DMA (Direct Memory Access) area used by the mlx4 driver. Consequently, a DMA leak occurred after removing a network device that used the driver. This update ensures that the mlx4 driver unmaps the correct portion of memory. As a result, the memory is freed correctly and no DMA leak occurs.

BZ#787771

Previously, when a memory allocation failure occurred, the mlx4 driver did not free the previously allocated memory correctly. Consequently, hotplug removal of devices using the mlx4 driver could not be performed. With this update, a memory allocation failure still occurs when the device MTU (Maximal Transfer Unit) is set to 9000, but hotplug removal the device is possible afer the failure.

BZ#759613

Due to a regression, the updated vmxnet3 driver used the ndo_set_features() method instead of various methods of the ethtool utility. Consequently, it was not possible to make changes to

vmxnet3-based network adapters in Red Hat Enterprise Linux 6.2. This update restores the ability of the driver to properly set features, such as csum or TSO (TCP Segmentation Offload), via ethtool.

BZ#759318

Previously, when a MegaRAID 9265/9285 or 9360/9380 controller got a timeout in the megaraid_sas driver, the invalid SCp.ptr pointer could be called from the megasas_reset_timer() function. As a consequence, a kernel panic could occur. An upstream patch has been provided to address this issue and the pointer is now always set correctly.

BZ#790673

The vmxnet3 driver in Red Hat Enterprise Linux 6.2 introduced a regression. Due to an optimization, in which at least 54 bytes of a frame were copied to a contiguous buffer, shorter frames were dropped as the frame did not have 54 bytes available to copy. With this update, transfer size for a buffer is limited to 54 bytes or the frame size, whichever is smaller, and short frames are no longer dropped in the described scenario.

BZ#755885

Previously, when isolating pages for migration, the migration started at the start of a zone while the free scanner started at the end of the zone. Migration avoids entering a new zone by never going beyond what the free scanner scanned. In very rare cases, nodes overlapped and the migration isolated pages without the LRU lock held, which triggered errors in reclaim or during page freeing. With this update, the isolate_migratepages() function makes a check to ensure that it never isolates pages from a zone it does not hold the LRU lock for, thus fixing this bug.

BZ#755380

Due to regression, an attempt to open a directory that did not have a cached dentry failed and the EISDIR error code was returned. The same operation succeeded if a cached dentry existed. This update modifies the nfs_atomic_lookup() function to allow fallbacks to normal look-up in the described scenario.

BZ#754356

Due to a race condition, the mac80211 framework could deauthenticate with an access point (AP) while still scheduling authentication retries with the same AP. If such an authentication attempt timed out, a warning message was returned to kernel log files. With this update, when deauthenticating, pending authentication retry attempts are checked and cancelled if found, thus fixing this bug.

BZ#692767

Index allocation in the virtio-blk module was based on a monotonically increasing variable "index". Consequently, released indexes were not reused and after a period of time, no new ones were available. Now, virtio-blk uses the ida API to allocate indexes, thus preventing this bug.

BZ#795441

When expired user credentials were used in the RENEW() calls, the calls failed. Consequently, all access to the NFS share on the client became unresponsive. With this update, the machine credentials are used with these calls instead, thus preventing this bug most of the time. If no machine credentials are available, user credentials are used as before.

BZ#753301

Previously, an unnecessary assertion could trigger depending on the value of the xpt_pool field. As a consequence, a node could terminate unexpectedly. The xpt_pool field was in fact unnecessary and this update removes it from the sunrpc code, thus preventing this bug.

BZ#753237

Prior to this update, the align_va_addr kernel parameter was ignored if secondary CPUs were initialized. This happened because the parameter settings were overridden during the initialization of secondary CPUs. Also, the align_va_addr parameter documentation contained incorrect parameter arguments. With this update, the underlying code has been modified to prevent the overriding and the documentation has been updated. This update also removes the unused code introduced by the patch for BZ#739456.

BZ#796277

Concurrent look-up operations of the same inode that was not in the per-AG (Allocation Group) inode cache caused a race condition, triggering warning messages to be returned in the unlock_new_inode() function. Although this bug could only be exposed by NFS or the xfsdump utility, it could lead to inode corruption, inode list corruption, or other related problems. With this update, the XFS_INEW flag is set before inserting the inode into the radix tree. Now, any concurrent look-up operation finds the new inode with XFS_INEW set and the operation is then forced to wait until XFS_INEW is removed, thus fixing this bug.

BZ#753030

Socket callbacks use the svc_xprt_enqueue() function to add sockets to the pool->sp_sockets list. In normal operation, a server thread will later take the socket off that list. Previously, on the nfsd daemon shutdown, still-running svc_xprt_enqueue() could re-add an socket to the sp_sockets list just before it was deleted. Consequently, system could terminate unexpectedly by memory corruption in the sunrpc module. With this update, the XPT_BUSY flag is put on every socket and svc_xprt_enqueue() now checks this flag, thus preventing this bug.

BZ#816034

Red Hat Enterprise Virtualization Hypervisor became unresponsive and failed to shut down or restart with the following message:

```
Shutting down interface breth0
```

This happened after configuring the NetConsole functionality with no bridge on top of a bond due to a mistake in the linking to the device structure. With this update, the linking has been fixed and the device binding is processed correctly in this scenario.

BZ#752528

When the md_raid1_unplug_device() function was called while holding a spinlock, under certain device failure conditions, it was possible for the lock to be requested again, deeper in the call chain, and causing a deadlock. With this update, md_raid1_unplug_device() is no longer called while holding a spinlock, thus fixing this bug.

BZ#797731

Previously, a bonding device had always the UFO (UDP Fragmentation Offload) feature enabled even when no slave interfaces supported UFO. Consequently, the tracepath command could not return correct path MTU. With this update, UFO is no longer configured for bonding interfaces by default if the underlying hardware does not support it, thus fixing this bug.

BZ#703555

When trying to send a kdump file to a remote system via the tg3 driver, the tg3 NIC (network interface controller) could not establish the connection and the file could not be sent. The kdump kernel leaves the MSI-X interrupts enabled as set by the crashed kernel, however, the kdump kernel

only enables one CPU and this could cause the interrupt delivery to the tg3 driver to fail. With this update, tg3 enables only a single MSI-X interrupt in the kdump kernel to match the overall environment, thus preventing this bug.

BZ#751087

On a system with an idle network interface card (NIC) controlled by the e1000e driver, when the card transmitted up to four descriptors, which delayed the write-back and nothing else, the run of the watchdog driver about two seconds later forced a check for a transmit hang in the hardware, which found the old entry in the TX ring. Consequently, a false "Detected Hardware Unit Hang" message was issued to the log. With this update, when the hang is detected, the descriptor is flushed and the hang check is run again, which fixes this bug.

BZ#750237

Previously, the idle_balance() function dropped or retook the rq->lock parameter, leaving the previous task vulnerable to the set_tsk_need_resched() function. Now, the parameter is cleared in setup_thread_stack() after a return from balancing and no successfully descheduled or never scheduled task has it set, thus fixing this bug.

BZ#750166

Previously, the doorbell register was being unconditionally swapped. If the Blue Frame option was enabled, the register was incorrectly written to the descriptor in the little endian format. Consequently, certain adapters could not communicate over a configured IP address. With this update, the doorbell register is not swapped unconditionally, rather, it is always converted to big endian before it is written to the descriptor, thus fixing this bug.

BZ#705698

The CFQ (Completely Fair Queuing) scheduler does idling on sequential processes. With changes to the IOeventFD feature, traffic pattern at CFQ changed and CFQ considered everything a thread was doing sequential I/O operations. Consequently, CFQ did not allow preemption across threads in Qemu. This update increases the preemption threshold and the idling is now limited in the described scenario without the loss of throughput.

BZ#798984

When short audio periods were configured, the ALSA PCM midlevel code, shared by all sound cards, could cause audio glitches and other problems. This update adds a time check for double acknowledged interrupts and improves stability of the snd-aloop kernel module, thus fixing this bug.

BZ#748559

Previously, the utime and stime values in the /proc/<pid>/stat file of a multi-threaded process could wrongly decrease when one of its threads exited. A backported patch has been provided to maintain monotonicity of utime and stime in the described scenario, thus fixing this bug.

BZ#800555

During tests with active I/O on 256 LUNs (logical unit numbers) over FCoE, a large number SCSI mid layer error messages were returned. As a consequence, the system became unresponsive. This bug has been fixed by limiting the source of the error messages and the hangs no longer occur in the described scenario.

BZ#714902

Previously, the compaction code assumed that memory on all cluster nodes is aligned to the same page-block size when isolating a cluster node for migration. However, when running a cluster on IBM

System x3850 X5 machines with two MAX 5 memory expansion drawers, memory is not properly aligned. Therefore, the isolate_migratepages() function could pass an invalid Page Frame Number (PFN) to the pfn_to_page() function, which resulted in a kernel panic. With this update, the compaction code has been modified so that the isolate_migratepages() function now calls the pfn_valid function to validate PFN when necessary, and the kernel no longer panics in this scenario described.

BZ#801730

The ctx->vif identifier is dereferenced in different parts of the iwlwifi code. When it was set to null before requesting hardware reset, the kernel could terminate unexpectedly. An upstream patch has been provided to address this issue and the crashes no longer occur in the described scenario.

BZ#717179

Previously, a CPU could service the idle load balancer kick request from another CPU, even without receiving the IPI. Consequently, multiple __smp_call_function_single() calls were done on the same call_single_data structure, leading to a deadlock. To kick a CPU, the scheduler already has the reschedule vector reserved. Now, the kick_process mechanism is used instead of using the generic smp_call_function mechanism to kick off the nohz idle load balancing and avoid the deadlock.

BZ#746484

A software bug related to Context Caching existed in the Intel IOMMU support module. On some newer Intel systems, the Context Cache mode has changed from previous hardware versions, potentially exposing a Context coherency race. The bug was exposed when performing a series of hot plug and unplug operations of a Virtual Function network device which was immediately configured into the network stack, i.e., successfully performed dynamic host configuration protocol (DHCP). When the coherency race occurred, the assigned device would not work properly in the guest virtual machine. With this update, the Context coherency is corrected and the race and potentially resulting device assignment failure no longer occurs.

BZ#746169

Due to a running cursor blink timer, when attempting to hibernate certain types of laptops, the i915 kernel driver could corrupt memory. Consequently, the kernel could crash unexpectedly. An upstream patch has been provided to make the i915 kernel driver use the correct console suspend API and the hibernate function now works as expected.

BZ#720611

Previously, the eth_type_trans() function was called with the VLAN device type set. If a VLAN device contained a MAC address different from the original device, an incorrect packet type was assigned to the host. Consequently, if the VLAN devices were set up on a bonding interface in Adaptive Load Balancing (ALB) mode, the TCP connection could not be established. With this update, the eth_type_trans() function is called with the original device, ensuring that the connection is established as expected.

BZ#806081

The slave member of "struct aggregator" does not necessarily point to a slave which is part of the aggregator. It points to the slave structure containing the aggregator structure, while completely different slaves (or no slaves at all) may be part of the aggregator. Due to a regression, the agg_device_up() function wrongly used agg->slave to find the state of the aggregator. Consequently, wrong active aggregator was reported to the /proc/net/bonding/bond0 file. With this update, agg->lag_ports->slave is used in the described scenario instead, thus fixing this bug.

BZ#806119

Due to the netdevice handler for FCoE (Fibre Channel over Ethernet) and the exit path blocking the keventd work queue, the destroy operation on an NPIV (N_Port ID Virtualization) FCoE port led to a deadlock interdependency and caused the system to become unresponsive. With this update, the destroy_work item has been moved to its own work queue and is now executed in the context of the user space process requesting the destroy, thus preventing this bug.

BZ#739811

Previously, when pages were being migrated via NFS with an active requests on them, if a particular inode ended up deleted, then the VFS called the truncate_inode_pages() function. That function tried to take the page lock, but it was already locked when migrate_page() was called. As a consequence, a deadlock occurred in the code. This bug has been fixed and the migration request is now refused if the PagePrivate parameter is already set, indicating that the page is already associated with an active read or write request.

BZ#808487

Previously, requests for large data blocks with the ZSECSENDCPRB ioctl() system call failed due to an invalid parameter. A misleading error code was returned, concealing the real problem. With this update, the parameter for the ZSECSENDCPRB request code constant is validated with the correct maximum value. Now, if the parameter length is not valid, the EINVAL error code is returned, thus fixing this bug.

BZ#809928

Due to incorrect use of the list_for_each_entry_safe() macro, the enumeration of remote procedure calls (RPCs) priority wait queue tasks stored in the tk_wait.links list failed. As a consequence, the rpc_wake_up() and rpc_wake_up_status() functions failed to wake up all tasks. This caused the system to become unresponsive and could significantly decrease system performance. Now, the list_for_each_entry_safe() macro is no longer used in rpc_wake_up(), ensuring reasonable system performance.

BZ#812259

Various problems were discovered in the iwlwifi driver happening in the 5 GHz band. Consequently, roaming between access points (AP) on 2.4 GHz and 5 GHz did not work properly. This update adds a new option to the driver that disables the 5 GHz band support.

BZ#810299

Previously, secondary, tertiary, and other IP addresses added to bond interfaces could overwrite the bond->master_ip and vlan_ip values. Consequently, a wrong IP address could be occasionally used, the MII (Media Independent Interface) status of the backup slave interface went down, and the bonding master interfaces were switching. This update removes the master_ip and vlan_ip elements from the bonding and vlan_entry structures, respectively. Instead, devices are directly queried for the optimal source IP address for ARP requests, thus fixing this bug.

BZ#727700

An anomaly in the memory map created by the mbind() function caused a segmentation fault in Hotspot Java Virtual Machines with the NUMA-aware Parallel Scavenge garbage collector. A backported upstream patch that fixes mbind() has been provided and the crashes no longer occur in the described scenario.

BZ#812108

Previously, with a transparent proxy configured and under high load, the kernel could start to drop packets, return error messages such as "ip_rt_bug: addr1 -> addr2, ?", and, under rare circumstances, terminate unexpectedly. This update provides patches addressing these issues and the described

problems no longer occur.

BZ#811815

The kdump utility does not support Xen para-virtualized (PV) drivers on Hardware Virtualized Machine (HVM) guests in Red Hat Enterprise Linux 6. Therefore, kdump failed to start if the guest had loaded PV drivers. This update modifies underlying code to allow kdump to start without PV drivers on HVM guests configured with PV drivers.

BZ#735105

When running a userspace program, such as the Ceph client, on the ext4 file system, a race condition between the sync/flush thread and the xattr-set thread could occur. This was caused by an incorrectly-set state flag on an inode. As a consequence, memory for the file system was incorrectly allocated, which resulted in file system corruption. With this update, the ext4 code has been modified to prevent this race condition from occurring and file systems are no longer corrupted under these circumstances.

BZ#728852

An unwanted interrupt was generated when a PCI driver switched the interrupt mechanism from the Message Signaled Interrupt (MSI or MSI-X) to the INTx emulation while shutting down a device. Due to this, an interrupt handler was called repeatedly, and the system became unresponsive. On certain systems, the interrupt handler of Intelligent Platform Management Interface (IPMI) was called while shutting down a device on the way to reboot the system after running kdump. In such a case, soft lockups were performed repeatedly and the shutdown process never finished. With this update, the user can choose not to use MSI or MSI-X for the PCI Express Native Hotplug driver. The switching between the interrupt mechanisms is no longer performed so that the unwanted interrupt is not generated.

BZ#731917

The time-out period in the qla2x00_fw_ready() function was hard-coded to 20 seconds. This period was too short for new QLogic host bus adapters (HBAs) for Fibre Channel over Ethernet (FCoE). Consequently, some logical unit numbers (LUNs) were missing after a reboot. With this update, the time-out period has been set to 60 seconds so that the modprobe utility is able to recheck the driver module, thus fixing this bug.

BZ#730045

Previously, the idmapper utility pre-allocated space for all user and group names on an NFS client in advance. Consequently, page allocation failure could occur, preventing a proper mount of a directory. With this update, the allocation of the names is done dynamically when needed, the size of the allocation table is now greatly reduced, and the allocation failures no longer occur.

BZ#811703

As part of mapping the application's memory, a buffer to hold page pointers is allocated and the count of mapped pages is stored in the do_dio field. A non-zero do_dio marks that direct I/O is in use. However, do_dio is only one byte in size. Previously, mapping 256 pages overflowed do_dio and caused it to be set to 0. As a consequence, when large enough number of read or write requests were sent using the st driver's direct I/O path, a memory leak could occur in the driver. This update increases the size of do_dio, thus preventing this bug.

BZ#728315

In the hpet_next_event() function, an interrupt could have occurred between the read and write of the HPET (High Performance Event Timer) and the value of HPET_COUNTER was then beyond that being written to the comparator (HPET_Tn_CMP). Consequently, the timers were overdue for up to

several minutes. Now, a comparison is performed between the value of the counter and the comparator in the HPET code. If the counter is beyond the comparator, the "-ETIME" error code is returned, which fixes this bug.

BZ#722297

In a Boot-from-San (BFS) installation via certain iSCSI adapters, driver exported sendtarget entries in the sysfs file system but the iscsistart failed to perform discovery. Consequently, a kernel panic occurred during the first boot sequence. With this update, the driver performs the discovery instead, thus preventing this bug.

BZ#805519

The SCSI layer was not using a large enough buffer to properly read the entire 'BLOCK LIMITS VPD' page that is advertised by a storage array. Consequently, the 'WRITE SAME MAX LEN' parameter was read incorrectly and this could result in the block layer issuing discard requests that were too large for the storage array to handle. This update increases the size of the buffer that the 'BLOCK LIMITS VPD' page is read into and the discard requests are now issued with proper size, thus fixing this bug.

BZ#803378

The Intelligent Platform Management Interface (IPMI) specification requires a minimum communication timeout of five seconds. Previously, the kernel incorrectly used a timeout of 1 second. This could result in failures to communicate with Baseboard Management Controllers (BMC) under certain circumstances. With this update, the timeout has been increased to five seconds to prevent such problems.

BZ#758404

The dm_mirror module can send discard requests. However, the dm_io interface did not support discard requests, and running an LVM mirror over a discard-enabled device led to a kernel panic. This update adds support for the discard requests to the dm_io interface, so that kernel panics no longer occur in the described scenario.

BZ#766051

Previously, when the schedule() function was run shortly after a boot, the following warning message was sometimes returned once per boot on the console:

```
5915: WARN_ON_ONCE(test_tsk_need_resched(next));
```

An upstream patch has been provided to address this issue and the WARN_ON_ONCE() call is no longer present in schedule(), thus fixing this bug.

BZ#786996

Prior to this update, bugs in the close() and send() functions caused delays and operation of these two functions took too long to complete. This update adds the IUCV_CLOSED state change and improves locking for close(). Also, the net_device handling has been improved in send(). As a result, the delays no longer occur.

BZ#770250

On NFS, when repeatedly reading a directory, content of which kept changing, the client issued the same readdir request twice. Consequently, the following warning messages were returned to the dmesg output:

> NFS: directory A/B/C contains a readdir loop.

This update fixes the bug by turning off the loop detection and letting the NFS client try to recover in the described scenario and the messages are no longer returned.

BZ#635817

A number of patches have been applied to the kernel in Red Hat Enterprise Linux 6.3 to improve overall performance and reduce boot time on extremely large UV systems (patches were tested on a system with 2048 cores and 16 TB of memory). Additionally, boot messages for the SGI UV2 platform were updated.

BZ#822697

Previously, if creation of an MFN (Machine Frame Number) was lazily deferred, the MFN could appear invalid when is was not. If at this point read_pmd_atomic() was called, which then called the paravirtualized __pmd() function, and returned zero, the kernel could terminate unexpectedly. With this update, the __pmd() call is avoided in the described scenario and the open-coded compound literal is returned instead, thus fixing this bug.

BZ#781566

Previously, on a system where intermediate P-states were disabled, the powernow-k8 driver could cause a kernel panic in the cpufreq subsystem. Additionally, not all available P-states were recognized by the driver. This update modifies the drive code so that it now properly recognizes all P-states and does not cause the panics in the described scenario.

BZ#783497

Due to an off-by-one bug in max_blocks checks, on the 64-bit PowerPC architecture, the tmpfs file system did not respect the size= parameter and consequently reported incorrect number of available blocks. A backported upstream patch has been provided to address this issue and tmpfs now respects the size= parameter as expected.

BZ#681906

This update introduces a performance enhancement which dramatically improves the time taken to read large directories from disk when accessing them sequentially. Large in this case means several hundred thousand entries or more. It does not affect the speed of looking up individual files (which is already fast), nor does it make any noticeable difference for smaller directories. Once a directory is cached, then again no difference can be noticed in performance. The initial read however, should be faster due to the readahead which this update introduces.

BZ#729586

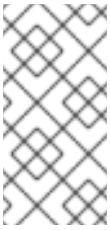
Red Hat Enterprise Linux 6.1 introduced naming scheme adjustments for emulated SCSI disks used with paravirtual drivers to prevent namespace clashes between emulated IDE and emulated SCSI disks. Both emulated disk types use the paravirt block device **xvd**. Consider the example below:

Table 5.1. The naming scheme example

	Red Hat Enterprise Linux 6.0	Red Hat Enterprise Linux 6.1 or later
emulated IDE	hda -> xvda	unchanged

	Red Hat Enterprise Linux 6.0	Red Hat Enterprise Linux 6.1 or later
emulated SCSI	sda -> xvda	sda -> xvde, sdb -> xvdf, ...

This update introduces a new module parameter, **xen_blkfront.sda_is_xvda**, that provides a seamless upgrade path from 6.0 to 6.3 kernel release. The default value of **xen_blkfront.sda_is_xvda** is **0** and it keeps the naming scheme consistent with 6.1 and later releases. When **xen_blkfront.sda_is_xvda** is set to **1**, the naming scheme reverts to the 6.0-compatible mode.



NOTE

Note that when upgrading from 6.0 to 6.3 release, if a virtual machine specifies emulated SCSI devices and utilizes paravirtual drivers and uses explicit disk names such as **xvd[a-d]**, it is advised to add the **xen_blkfront.sda_is_xvda=1** parameter to the kernel command line before performing the upgrade.

BZ#756307

In previous Red Hat Enterprise Linux 6 releases, the kernel option xen_emul_unplug=never did not disable xen platform pci device and that lead to using para-virtual devices instead of emulated ones. This fix, in addition to fixing the irq allocation issue for emulated network devices, allows to disable para-virtual drivers using the xen_emul_unplug=never kernel option as described in "Virtualization Guide: Edition 5.8" chapter "12.3.5. Xen Para-virtualized Drivers on Red Hat Enterprise Linux 6".

BZ#749251

When a process isolation mechanism such as LXC (Linux Containers) was used and the user space was running without the CAP_SYS_ADMIN identifier set, a jailed root user could bypass the dmesg_restrict protection, creating an inconsistency. Now, writing to dmesg_restrict is only allowed when the root has CAP_SYS_ADMIN set, thus preventing this bug.

BZ#788591

Previously, the code for loading multipath tables attempted to load the scsi_dh module even when it was already loaded, which caused the system to become unresponsive. With this update, the code does not attempt to load the scsi_dh module when it is already loaded and multipath tables are loaded successfully.

BZ#801877

Due to an error in the code for ASPM (Active State Power Management) tracking, the system terminated unexpectedly after attempts to remove a PCI bus with both PCIe and PCI devices connected to it when PCIe ASPM was disabled using the "pcie_aspm=off" kernel parameter. This update ensures that the ASPM handling code is not executed when ASPM is disabled and the server no longer crashes in the aforementioned scenario.

BZ#804608

Due to an error in the underlying source code, the perf performance counter subsystem calculated event frequencies incorrectly. This update fixes the code and calculation of event frequencies now returns correct results.

BZ#787771

Previously, when a memory allocation failure occurred, the mlx4 driver did not free the previously allocated memory correctly. Consequently, hotplug removal of devices using the mlx4 driver could not be performed. With this update, a memory allocation failure still occurs when the device MTU (Maximal Transfer Unit) is set to 9000, but hotplug removal the device is possible afer the failure.

BZ#787762

Previously, an incorrect portion of memory was freed when unmapping a DMA (Direct Memory Access) area used by the mlx4 driver. Consequently, a DMA leak occurred after removing a network device that used the driver. This update ensures that the mlx4 driver unmaps the correct portion of memory. As a result, the memory is freed correctly and no DMA leak occurs.

BZ#812415

The Intel SCU driver did not properly interact with the system BIOS to honor the Spread Spectrum Clock (SSC) settings and state by the BIOS controls: even though the SSC mode was enabled in the preboot BIOS environment, it became disabled after boot due to incorrect parameter parsing from the ROM option. With this update, the kernel driver has been modified to correctly parse OEM parameters from the ROM option and the problem no longer occurs.

BZ#811023

The iw_cxgb4 driver has been updated so as to fix a race that occurred when an ingress abort failed to wake up the thread blocked in rdma_init() causing the application to become unresponsive. Also, the driver has been modified to return and not to call the wake_up() function if no endpoint is found as this is not necessary.

BZ#818371

When creating a snapshot of a mounted RAID volume, a kernel panic could occur. This happened because a timer designed to wake up an I/O processing thread was not deactivated when the RAID device was replace by a snapshot origin. The timer then woke a thread that attempted to access memory that had already been freed, resulting in a kernel panic. With this update, this bug has been fixed and the kernel panic no longer occurs in this scenario.

BZ#821329

Previously, attempts to add a write-intent bitmap to an MD array using v1.0 metadata and then using the array without rebooting caused a kernel OOPS. This occurred because the kernel did not reload the bitmap information correctly after creating the bitmap. With this update, the kernel loads the information correctly on bitmap creation, as expected and the kernel OOPS no longer occurs.

BZ#817090

On IBM System z, a kernel panic could occur if there was high traffic workload on HiperSockets devices. This happened due to a conflict in the qeth driver between asynchronous delivery of storage blocks for HiperSockets devices and outdated SIGA (System Information GAthering) retry code. With this update, the SIGA retry code has been removed from the qeth driver and the problem no longer occurs.

BZ#736931

Previously, certain internal functions in the real-time scheduler only iterated over runnable real-time tasks instead of iterating over all existing tasks. Consequently, when processing multiple real-time threads on multiple logical CPUs and one CPU was disabled, the kernel could panic with the following error message:

```
kernel BUG at kernel/sched_rt.c:460!
```

This update modifies the real-time scheduler so that all real-time tasks are processed as expected and the kernel no longer crashes in this scenario.

BZ#756301

Due to a bug in the qla2xxx driver and the HBA firmware, storage I/O traffic could become unresponsive during storage fault testing. With this update, these bugs have been fixed and storage traffic no longer hangs in the described scenario.

BZ#767505

When resetting a virtual block device and a config interrupt was received, the config_work handler could attempt to access the device configuration after the device had already been removed from the system but before the device was reset. This resulted in a kernel panic. With this update, the underlying code has been modified to use a mutex lock and disable the device configuration during the reset. Config interrupts can no longer be processed during the reset of the virtual block device and the kernel no longer panics in this scenario.

BZ#784430

After some recent changes in USB driver code, previous versions of the kernel did not handle, under some circumstances, standard and warm reset of USB3.0 ports correctly. Consequently, the system was not able to detect and automatically mount a USB3.0 device when the device was re-attached to a USB3.0 port after it was unmounted. This update applies several upstream patches related to handling USB3.0 ports, and USB3.0 devices are now automatically re-attached as expected in the scenario described.

BZ#738491

Previously, the mlx4 driver expected Remote Direct Memory Access (RDMA) communication to be performed over an InfiniBand link layer and the driver thus used the InfiniBand link layer part of the code to record transfer statistics. However, Mellanox RDMA over Converged Ethernet (RoCE) devices use an Ethernet link layer for RDMA communication, which caused that RDMA communication was not accounted under these circumstances, and the displayed statistics were incorrect. With this update, the underlying code has been modified so that the driver now uses a "global" counter for RDMA traffic accounting on Ethernet ports, and users can see correct RDMA transfer statistics.

BZ#749059

Due to a missing validation check, the mlx4 driver could attempt to access an already freed data element in the core network device structure of the network layer. As a consequence, if a Mellanox ConnectX HCA InfiniBand adapter was unexpectedly removed from the system while the adapter processed ongoing Remote Direct Memory Access (RDMA) communication, the kernel panicked. With this update, the mlx4 driver has been modified to verify that the core network device structure is valid before attempting to use it for outgoing communication. The kernel now no longer panics when an adapter port is unexpectedly disabled.

Enhancements



NOTE

For more information on the most important of the RHEL 6.3 kernel enhancements, refer to the *Kernel* and *Device Drivers* chapters in the Red Hat Enterprise Linux 6.3 Release Notes.

For a summary of added or updated **procfs** entries, **sysfs** default values, boot parameters, kernel configuration options, or any noticeable behavior changes, refer to Chapter 1, *Important Changes to External Kernel Parameters* .

BZ#808315

LED support has been added to the sysfs interfaces.

BZ#805658

The WinFast VP200 H (Teradici) snd-hda-intel audio device has been added, and is recognized by the alsa driver.

BZ#744301

The Brocade BFA Fibre Channel and FCoE driver is no longer a Technology Preview. In Red Hat Enterprise Linux 6.3 the BFA driver is fully supported.

BZ#744302

The Brocade BNA driver for Brocade 10Gb PCIe ethernet Controllers is no longer a Technology Preview. In Red Hat Enterprise Linux 6.3 the BNA driver is fully supported.

BZ#696383

Persistent storage (pstore), a file system interface for platform dependent persistent storage, now supports UEFI.

BZ#661765

This release adds support for a new kernel auditing feature that allows for inter-field comparisons. For each audit event, the kernel collects information about what is causing the event. Now, you can use the "-C" command to tell the kernel to compare: auid, uid, euid, suid, fsuid, or obj_uid; and gid, egid, sgid, fsgid, or obj_gid. The two groups cannot be mixed. Comparisons can use either of the equal or not equal operators.

BZ#821561

This update adds the rh_check_unsupported() function and blacklists unsupported future Intel processors.

BZ#786997

When AF_IUCV sockets were using the HiperSockets transport, maximum message size for such transports depended on the MTU (maximum transmission unit) size of the HiperSockets device bound to a AF_IUCV socket. However, a socket program could not determine maximum size of a message. This update adds the MSGSIZE option for the getsockopt() function. Through this option, the maximum message size can be read and properly handled by AF_IUCV.

BZ#596419

The cred argument has been included in the security_capable() function so that it can be used in a wider range of call sites.

BZ#773052

Red Hat Enterprise Linux 6.3 adds support for the Wacom Cintiq 24HD (a 24-inch Drawing Tablet).

BZ#738720

This update adds additional fixed tracepoints to trace signal events.

BZ#704003

This update adds the missing raid6test.ko module.

BZ#788634

The keyrings kernel facility has been upgraded to the upstream version, which provides a number of bug fixes and enhancements over the previous version. In particular, the garbage collection mechanism has been re-worked.

BZ#788156

The perf tool has been upgraded to upstream version 3.3-rc1, which provides a number of bug fixes and enhancements over the previous version.

BZ#766952

The wireless LAN subsystem has been updated. It introduces the dma_unmap state API and adds a new kernel header file: include/linux/pci-dma.h.

BZ#723018

The dm-thinp targets, thin and thin-pool, provide a device mapper device with thin-provisioning and scalable snapshot capabilities. This feature is available as a Technology Preview.

BZ#768460

In Red Hat Enterprise Linux 6.3, SHA384 and SHA512 HMAC authentication algorithms have been added to XFRM.

Users should upgrade to these updated packages, which contain backported patches to correct these issues, fix these bugs, and add these enhancement. The system must be rebooted for this update to take effect.

5.135.15. RHSA-2013:0662 – Important: kernel security and bug fix update

Updated kernel packages that fix one security issue and several bugs are now available for Red Hat Enterprise Linux 6.3 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 Fix

CVE-2013-0871, Important

This update fixes the following security issue:

* A race condition was found in the way the Linux kernel's ptrace implementation handled PTRACE_SETREGS requests when the debuggee was woken due to a SIGKILL signal instead of being stopped. A local, unprivileged user could use this flaw to escalate their privileges.

Bug Fixes

BZ#908735

Previously, init scripts were unable to set the MAC address of the master interface 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 setting the MAC address of the first slave interface as the MAC address of the master interface.

BZ#909158

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#915582

Due to the incorrect validation of a pointer dereference in the d_validate() function, running a command such as ls or find on the MultiVersion File System (MVFS), used by IBM Rational ClearCase, for example, could trigger a kernel panic. This update modifies d_validate() to verify the parent-child dentry relationship by searching through the parent's d_child list. The kernel no longer panics in this situation.

BZ#916956

A previously backported patch introduced usage of the page_descs length field but did not set the page data length for the FUSE page descriptor. This code path can be exercised by a loopback device (pagecache_write_end) if used over FUSE. As a result, fuse_copy_page does not copy page data from the page descriptor to the user-space request buffer and the user space can see uninitialized data. This could previously lead to file system data corruption. This problem has been fixed by setting the page_descs length prior to submitting the requests, and FUSE therefore provides correctly initialized data.

Users should upgrade to these updated packages, which contain backported patches to resolve these issues. The system must be rebooted for this update to take effect.

5.135.16. RHSA-2013:0832 – Important: kernel security update

Updated kernel packages that fix one security issue are now available for Red Hat Enterprise Linux 6.3 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 Fix

CVE-2013-2094, Important

This update fixes the following security issue:

* It was found that the Red Hat Enterprise Linux 6.1 kernel update (RHSA-2011:0542) introduced an integer conversion issue in the Linux kernel's Performance Events implementation. This led to a user-supplied index into the perf_swevent_enabled array not being validated properly, resulting in out-of-bounds kernel memory access. A local, unprivileged user could use this flaw to escalate their privileges.

A public exploit that affects Red Hat Enterprise Linux 6 is available.

Refer to Red Hat Knowledge Solution 373743, linked to in the References, for further information and mitigation instructions for users who are unable to immediately apply this update.

Users should upgrade to these updated packages, which contain a backported patch to correct this issue. The system must be rebooted for this update to take effect.

5.135.17. RHSA-2013:1450 – 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 6.3 Extended Update Support.

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-2013-2224, Important

It was found that the fix for CVE-2012-3552 released via RHSA-2012:1540 introduced an invalid free flaw in the Linux kernel's TCP/IP protocol suite implementation. A local, unprivileged user could use this flaw to corrupt kernel memory via crafted sendmsg() calls, allowing them to cause a denial of service or, potentially, escalate their privileges on the system.

CVE-2013-4299, Moderate

An information leak flaw was found in the way Linux kernel's device mapper subsystem, under certain conditions, interpreted data written to snapshot block devices. An attacker could use this flaw to read data from disk blocks in free space, which are normally inaccessible.

CVE-2013-2852, Low

A format string flaw was found in the b43_do_request_fw() function in the Linux kernel's b43 driver implementation. A local user who is able to specify the "fwpostfix" b43 module parameter could use this flaw to cause a denial of service or, potentially, escalate their privileges.

Red Hat would like to thank Fujitsu for reporting CVE-2013-4299, and Kees Cook for reporting CVE-2013-2852.

Bug Fixes

BZ#1004185

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#1007467

A race condition in the abort task and SPP device task management path of the isci 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 isci 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#1008507

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.

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.

5.135.18. RHBA-2013:1190 – kernel bug fix update

Updated kernel packages that fix several bugs are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The kernel packages contain the Linux kernel, the core of any Linux operating system.

Bug Fixes

BZ#979291

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#982114

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#982469

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#988226

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#990087

BE family hardware could falsely indicate an unrecoverable error (UE) on certain platforms and stop further access to be2net-based network interface cards (NICs). A patch has been applied to disable the code that stops further access to hardware for BE family network interface cards (NICs). For a real UE, it is not necessary as the corresponding hardware block is not accessible in this situation.

BZ#991344

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.

Users should upgrade to these updated packages, which contain backported patches to correct these bugs. The system must be rebooted for this update to take effect

5.136. KEXEC-TOOLS

5.136.1. RHBA-2012:1554 – kexec-tools bug fix update

Updated kexec-tools packages that fix one bug 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 Fix

BZ#878371

Previously, the kexec utility incorrectly recognized the Xen DomU (HVM) guest as the Xen Dom0 management domain. Consequently, the kernel terminated unexpectedly and the kdump utility generated the vmcore file with no NT_PRSTATUS notes. The crash also led to NULL pointer dereference. With this update, kexec collects positions and sizes of NT_PRSTATUS from /sys/devices/system/cpu/cpuN/crash_notes on Xen DomU and from /proc/iomem on Xen Dom0. As a result, the crashes no longer occur.

All users of kexec-tools are advised to upgrade to these updated packages, which fix this bug.

5.136.2. RHBA-2012:0758 – kexec-tools bug fix update

Updated kexec-tools packages that fix multiple bugs and add several enhancements are now available for Red Hat Enterprise Linux 6.

The kexec-tools package provides the **/sbin/kexec** binary that facilitates a new kernel to boot using the kernel kexec feature. This package contains ancillary utilities that together with the binary form the user-space component of the kernel kexec feature.

Bug Fixes

BZ#821930

The **kdump** utility failed to find the member device if a member device of a bridge was renamed, that is, it was not using its default device name. Consequently, bridge mapping and kdump over network failed. With this update, the bridge member details are acquired from transformed **ifcfg** files when gaining the bridge member name for the dump kernel and kdump over network succeeds in this scenario.

BZ#798886

The **kdump** utility does not support Xen para-virtualized (PV) drivers on Hardware Virtualized Machine (HVM) guests. Therefore, kdump failed to start if the guest had PV drivers. This update modifies the code that allows kdump to start without PV drivers on HVM guests configured with PV drivers.

BZ#752458

Previously, **kdump** did not handle dumping over a bonding device in an IEEE 802.1Q network correctly and failed when requesting a core dump through such a device. With this update, code to allow kdump to handle VLAN tagging, so as to operate on such bonding devices properly, has been added and kdump succeeds under these circumstances.

BZ#812816

On IBM System z architectures, kdump request over network failed as no core dump was found on the remote server. This happened because network devices were not brought online before IP configuration. The IP configuration, therefore, referred to a non-existing network interface and the connection failed. With this update, network devices are brought online before the IP information is set.

BZ#805803

Previously, **kdump** did not bring up a bonding device and failed if the **BOOTPROTO** property in the **ifcfg-*<device>*** networking script of the bonding device was set to **none**. This happened because the **mkdumprd** utility did not handle the **BOOTPROTO** setting correctly. With this update, **mkdumprd** handles the setting correctly and **kdump** succeeds when dumping remotely over a bonding device with the **BOOTPROTO=none** setting.

BZ#802201

On x86 architectures, if the **crashkernel** boot option is set to **auto** to allow automatic reservation of memory for the kdump kernel, the threshold memory changes to 2 . However, the **firstboot** application was incorrectly using the 4 GB threshold. With this update, **firstboot** uses the same threshold value as the kernel.

BZ#785264

When kdump did not have permissions to dump to a NFS (Network File System) target, a kernel panic occurred. This happened because the init script of the kdump kernel did not check permissions to the target NFS resource after kdump had already started. With this update, the init script checks

the NFS directory permissions and kdump performs the default action specified by the user for situations when NFS dump fails due to lacking permissions.

BZ#738290

Previously, kdump failed if run with FCoE HBA Driver (fnic) and iSCSI dump targets. This update adds support for iSCSI targets using software initiators without iBFT. Note that iSCSI targets using hardware initiators are not supported.

BZ#814629

The **mkdumprd** utility uses the root device as the default dump target for the kdump initrd if the **/etc/kdump.conf** file does not define a dump target. However, mkdumprd used the device name instead of its UUID (Universally Unique Identifier), which could cause kdump to fail. With this update, the device UUID is used instead of its device name by default and kdump over root device succeeds in the scenario described.

BZ#743551

Previously, kdump could not capture the core dump when the target partition was encrypted and errors occurred. With this update, kdump warns the user when they specify an encrypted device as a dump target.

BZ#748654

Due to a bug in the **makedumpfile** utility, the utility failed when used to re-filter the vmcore core dump. With this update, **makedumpfile** has been modified to handle the input from vmcore correctly.

BZ#771671

Previously, when **makedumpfile** was redirected using the pipeline to **ssh** and failed, kdump did not drop the shell to the user even though the **default_action** property was set to **shell** in the **kdump.conf** file. With this update, the pipeline redirection fails as soon as makedumpfile fails and the shell is dropped immediately in the scenario described.

BZ#781919

Previously, due to a bug in the **mkdumprd** utility, data on an NFS server could be removed when the NFS unmount process failed. With this update, the problem has been fixed and the original data on the NFS server is now retained unchanged if unmounting fails.

BZ#761488

The kdump kernel did not clear the MCE (Machine Check Exception) status propagated from the kernel; the kdump kernel continued to boot without clearing the MCE status bits and triggered the MCE error again. With this update the MCE error in the kernel is not passed to the kdump kernel.

BZ#805464

The **kdump** utility did not bring DASDs (Direct Access Storage Devices) online before copying the vmcore file on IBM System z architectures. Consequently, kdump was waiting for the device and became unresponsive. With this update, DASD devices are brought online before copying vmcore and kdump works as expected in the scenario described.

BZ#753756

When running **kdump** after a kernel crash on a system with an **ext4** file system, the kdump initrd (initial RAM disk) could have been created with zero-byte size. This happened because the system waits for several seconds before writing the changes to the disk on an ext4 file system.

Consequently, the kdump initial root file system (rootfs) could not be mounted and kdump failed. This update modifies kexec-tools so that it perform the sync operations after creating initrd. This ensures that initrd is properly written to the disk before trying to mount rootfs, and kdump now successfully proceeds and captures the core dump on systems with an ext4 file system.

BZ#759003

When using SSH or NFS, it was not possible to capture the vmcore file if using a static IP without a gateway. This happened because the **mkdumprd** utility wrote an empty value as the gateway address into initrd. With this update, the gateway address is not assigned any value and kdump in such environments succeeds.

BZ#782674

On IBM System z, the **makedumpfile** command could fail because it did not translate virtual addresses to physical addresses correctly. With this update, **makedumpfile** handles virtual addresses correctly on these architectures and the command execution succeeds in this scenario.

BZ#784114

The **kdump** utility failed to load proper fonts for the kdump shell. Consequently, colored characters were returned in the Cyrillic alphabet in the kdump shell. With this update, the console fonts are installed in kdump initrd and the kdump default shell returns colored characters in the Latin alphabet as expected.

BZ#794580

The **mkdumprd** utility handled only the default path (the **/lib/modules/*<kernelVersion>*/** directory) of a modprobe and did not cover other module directories. Consequently, mkdumprd failed if there were modules located in other that the default path directory. The **mkdumprd** utility now handles **/lib/modules/*<kernelVersion>*/updates/** as well as the **/lib/modules/*<kernelVersion>*/** directory and mkdumprd succeeds under these circumstances.

BZ#697657

Previously, even though the SELinux policycoreutils package was not installed, mkdumprd used the **sestatus** and **setenforce** utilities. Consequently, kdump threw the following error while propagating ssh keys:

```
/etc/init.d/kdump: line 281: /usr/sbin/sestatus: No such file or directory
```

With this update, mkdumprd acquires information on the policycoreutils availability and processes the SELinux attribute using the **sysfs** tool.



NOTE

When the policycoreutils package is removed, SELinux must be disabled by adding the **selinux=0** option to the kernel command line. A system with SELinux enabled and the policycoreutils package not installed is considered a broken environment in which kdump returns the aforementioned errors. When you remove the policycoreutils package, make sure you have also disabled SELinux with **selinux=0**; otherwise, the problem will preserve.

BZ#801497

The restricted shell (**rksh**) does not allow redirections using a pipeline. Consequently, kdump failed if the remote user that was used when requesting the core dump was configured with a restricted shell.

With this update, the **dd** command is used instead of **cat** to copy vmcore, and kdump succeeds when a remote user uses the rksh shell.

BZ#635583

Previously, kdump could become unresponsive if called through a wireless interface as this option is not supported even though the **iwlwifi** (Intel Wireless WiFi Link) modules for interface devices were included in kdump initrd. With this update, the **iwlwifi** modules are no longer loaded into kdump initrd.

BZ#600575

Previously, the order in which kdump loaded storage drivers caused that a USB-attached storage was sometimes not correctly detected on certain 32-bit x86 systems. Consequently, devices were enumerated wrongly, and dumps therefore failed. The code has been fixed and the core dump takes place successfully.

BZ#699318

The **mkdumprd** utility ignored the **PREFIX** variable setting and the **ifconfig** utility failed during core dumping over network. With this update, the **mkdumprd** utility handles the **PREFIX** variable setting in **ifcfg-<device>** network scripts correctly.

BZ#794981

When the dump target was a raw device, the kdump init script created an unnecessary directory and an empty vmcore file in the **/var/crash/** directory. With this update, kdump checks the device header of the target device. If the header is invalid, kdump does not handle the situation as a crash and the redundant resources are no longer created on raw devices.

BZ#729675

When booting the **kdump** environment failed, kdump mounted the root device and ran the init script in user space if no default action was specified in the **kdump.conf** file. However, running init in user space to capture the core dump could cause an OOM (Out of Memory) state in the dump kernel. With this update, the kernel is now rebooted by default under these circumstances. Also, a new default option, **mount_root_run_init**, has been added to kdump. With this option, the kernel mounts the root partition, and runs the init and kdump service to try to save the kernel core dump, which allows the user to apply the previous behavior of kdump.

Enhancements

BZ#738866

Support for kdump on IBM System z has been added.

BZ#805040

The **firstboot** utility now supports configuring kdump for IBM S/390 architectures.

BZ#694498

Previously, the vmcore code dump could contain sensitive information, such as security keys, and potentially leak security information of the root user. With this update, the makedumpfile tool filters out such sensitive kernel data and vmcore no longer contains any sensitive security information.

BZ#738864

On IBM System z, the **makedumpfile** utility has been improved to handle vmcore correctly and the output no longer contains spurious errors.

BZ#795804

The **kdump** utility now supports the NFSv4 file system format.

BZ#736886

The **makedumpfile** can now handle Fujitsu's **sadump** dump format.

BZ#727413

The **kdump** utility now supports multipath storage devices as its dump targets.

Users of kexec-tools should upgrade to these updated packages, which fix these bugs and add these enhancements.

5.137. KEYUTILS

5.137.1. RHEA-2012:0963 – keyutils enhancement update

Updated keyutils packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

The keyutils package provides utilities to control the Linux kernel key management facility and to provide a mechanism by which the kernel calls up to user space to get a key instantiated.

Enhancement

BZ#772497

With this update, the request-key utility allows multiple configuration files to be provided. The request-key configuration file and its associated key-type specific variants are used by the request-key utility to determine which program should be run to instantiate a key.

All users of keyutils are advised to upgrade to these updated packages, which add this enhancement.

5.138. KRB5

5.138.1. RHBA-2012:1294 – krb5 bug fix update

Updated krb5 packages are now available for Red Hat Enterprise Linux 6.

Kerberos is a network authentication system which allows clients and servers to authenticate to each other using symmetric encryption and a trusted third-party, the Key Distribution Center (KDC).

Bug Fix

BZ#852455

Due to a previous update to a locally-applied patch, files created by the libkrb5 library were given correct SELinux labels. However, each flushing of the replay cache caused the file context configuration to be reloaded, to ensure that the correct label is applied to the newly-created replacement replay cache file. This resulted in large performance degradation in applications which accept authentication and use replay caches. With this update, the context configuration is only

loaded when the context configuration file has been modified and the configuration is now freed only when the library is unloaded or the calling application exits, thus greatly lowering the impact of this problem.

Users of krb5 are advised to upgrade to these updated packages, which fix this bug.

5.138.2. RHSA-2012:1131 – Important: krb5 security update

Updated krb5 packages that fix two security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

Kerberos is a network authentication system which allows clients and servers to authenticate to each other using symmetric encryption and a trusted third-party, the Key Distribution Center (KDC).

Security Fixes

CVE-2012-1015

An uninitialized pointer use flaw was found in the way the MIT Kerberos KDC handled initial authentication requests (AS-REQ). A remote, unauthenticated attacker could use this flaw to crash the KDC via a specially-crafted AS-REQ request.

CVE-2012-1013

A NULL pointer dereference flaw was found in the MIT Kerberos administration daemon, kadmind. A Kerberos administrator who has the "create" privilege could use this flaw to crash kadmind.

Red Hat would like to thank the MIT Kerberos project for reporting CVE-2012-1015. Upstream acknowledges Emmanuel Bouillon (NCI Agency) as the original reporter of CVE-2012-1015.

All krb5 users should upgrade to these updated packages, which contain backported patches to correct these issues. After installing the updated packages, the krb5kdc and kadmind daemons will be restarted automatically.

5.138.3. RHBA-2012:0921 – krb5 bug fix and enhancement update

Updated krb5 packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Kerberos is a network authentication system which allows clients and servers to authenticate to each other using symmetric encryption and a trusted third party, the Key Distribution Center (KDC).

Bug Fixes

BZ#748528

When obtaining initial credentials using a keytab file, a client failed to generate encrypted timestamp pre-authentication data to be sent to a KDC. This happened if the keytab file did not include a key of the first encryption type which the server suggested that the client could use, even if the client possessed other keys which would have been acceptable for this purpose. Attempts to use the kinit command with a keytab file often failed when the keytab file did not contain AES keys, but the

client's libraries and the KDC both supported AES. If the client libraries and the KDC both support at least one encryption type for which the keytab contains a key, the client now succeeds in obtaining credentials.

BZ#752405

If the KDC was started with the "-w" flag, and one of the worker processes which it started exited abnormally, the KDC failed to correctly update its count of the number of child processes which were still running. Consequently, the KDC waited for one or more of its worker processes to exit when it was shut down but because those processes had already exited it would never finish. This update backports a fix to correctly account for the number of running processes and KDC now shuts down correctly in the scenario described.

BZ#761006

If a GSS acceptor application exported its security context, the file handle for the replay cache which it had used while establishing the security context would not be properly closed. Consequently, the number of opened files increased until the limit for the process was reached. When this happened in rpc.svcgssd, which exports all of its contexts in order to pass information to the kernel, the daemon became unresponsive. This update backports the fix for this bug and the file handles are now properly closed.

BZ#813883

When the system was authenticated to a Windows AD using SSSD, the Kerberos credentials cache files created after login were mislabeled with an incorrect SELinux context. This was because the SELinux context was not re-created for a new replay cache, and instead the context of the old replay cache was used for new files. Kerberos credential cache files are now properly labelled with a correct SELinux context.

BZ#801033

When uninstalling the krb5-workstation package, info pages in the package were being removed from the info page index after the files were already removed. Info pages are now removed from the info page index before they are removed.

BZ#786216

When a client asks a KDC for a ticket, it can set a flag (the canonicalize bit) in its request indicating that it will accept a referral to another realm if that service is in a different realm. if the service is in a different realm, the KDC may then reply with a cross-realm TGT, indicating that the request should be made to a different realm. In come cases, for example when obtaining password-changing credentials, a referral TGT for the same realm was generated. This could create a loop in the process which was caught and an error was returned, failing to acquire the password-changing credentials. With this update, the same request is now retried without the canonicalize bit set, which elicits the desired result from the KDC.

Enhancements

BZ#761523

This update backports modifications to the Kerberos client which allows server applications to store credentials which have been obtained using s4u2proxy in a credential cache.

BZ#799161

This update backports modifications which allow GSS acceptor applications to better accept authentication from clients which use mechanisms other than the server's default, but which the server could still support.

BZ#782211

Previous versions of Red Hat Enterprise Linux contained modifications to allow Kerberos-aware services to accept authentication requests which were encrypted using keys marked with version number "0" regardless of the version number used in the keytab. While this is now the default behavior when a service does not specify its principal name to the APIs which it uses, the krb5_verify_init_creds() function and applications which use it, still required the modifications to support these cases. This update reintroduces them.

Users are advised to upgrade to these updated krb5 packages, which provide numerous bug fixes and enhancements.

5.139. KSH

5.139.1. RHBA-2012:0952 – ksh bug fix update

Updated ksh packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

KSH-93 is the most recent version of the KornShell by David Korn of AT&T Bell Laboratories. KornShell is a shell programming language which is also compatible with sh, the original Bourne Shell.

Bug Fixes

BZ#577223

Previously, ksh sometimes did not restore terminal settings after read timeout when operating in a multibyte environment. This could cause the terminal to no longer echo input characters. This updates applies a patch ensuring that the terminal is restored properly after the timeout and the user's input is now echoed as expected.

BZ#742930

When exiting a subshell after a command substitution, ksh could prematurely exit without any error message. With this update, ksh no longer terminates under these circumstances and all subsequent commands are processed correctly.

BZ#743840

Previously, ksh did not prevent modifications of variables of the read-only type. As a consequence, ksh terminated unexpectedly with a segmentation fault when such a variable was modified. With this update, modification of read-only variables are not allowed, and ksh prints an error message in this scenario.

BZ#781498

Previously, ksh did not close certain file descriptors prior to execution. This could lead to a file descriptor leak, and certain applications could consequently report error messages. With this update, file descriptors are marked to be closed on execution if appropriate, so file descriptor leaks no longer occur.

BZ#781976

In certain cases, ksh unnecessarily called the vfork() function. An extra process was created, and it could be difficult to determine how many instances of a script were running. A patch has been applied to address this problem, and extra processes are no longer created if not required.

BZ#786787

Previously, ksh could incorrectly seek in the input stream. This could lead to data corruption in the here-document section of a script. This update corrects the seek behavior, so the data no longer gets corrupted in this scenario.

BZ#798868

Previously, ksh did not allocate the correct amount of memory for its data structures containing information about file descriptors. When running a task that used file descriptors extensively, ksh terminated unexpectedly with a segmentation fault. With this update, the proper amount of memory is allocated, and ksh no longer crashes if file descriptors are used extensively.

BZ#800684

Previously, ksh did not expand the tilde (~) character properly. For example, characters in the tilde prefix were not treated as a login name but as a part of the path and the "No such file or directory" message was displayed. The underlying source code has been modified and tilde expansion now works as expected in such a scenario.

All users of ksh are advised to upgrade to these updated packages, which fix these bugs.

5.139.2. RHBA-2013:1019 – ksh bug fix update

Updated ksh packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

KornShell (ksh) is a Unix shell developed by Bell Labs, which is backward-compatible with the Bourne shell and includes many features of the C shell. KornShell complies with POSIX.2 [Shell and Utilities, Command Interpreter] standard.

Bug Fix

BZ#927584

Previously, the output a 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 now the here-document contains the output of command substitutions as expected.

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

5.140. LATENCYTOP

5.140.1. RHBA-2012:0864 – latencytop bug fix and enhancement update

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

LatencyTOP is a tool to monitor system latencies.

Bug Fix

BZ#633698

When running LatencyTOP as a normal user, LatencyTOP attempted and failed to mount the debug file system. A misleading error message was displayed, suggesting that kernel-debug be installed even though this was already the running kernel. LatencyTOP has been improved to exit and display "Permission denied" when run as a normal user. In addition, fsync view has been removed from the "latencytop" package because it depended on a non-standard kernel tracer that was never present in Red Hat Enterprise Linux kernels or upstream kernels. As a result, LatencyTOP no longer attempts to mount the debugfs file system.

Enhancement

BZ#726476

The "latencytop" package requires GTK libraries. Having GTK libraries installed on servers may be undesirable. A build of LatencyTOP without dependencies on GTK libraries is now available under the package name "latencytop-tui".

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

5.141. LIBBONOBO

5.141.1. RHBA-2012:0908 – libbonobo bug fix update

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

The libbonobo packages contain the libraries for bonobo, which is a component system based on CORBA, used by the GNOME desktop.

Bug Fix

BZ#728458

Prior to this update, the activation server handled parse errors incorrectly, for example as incorrectly nested tags, in server files. As a consequence, the activation server could, under certain circumstances, unexpectedly terminate with an abort signal upon encountering invalid server files in the activation path. This update modifies the underlying code so that invalid files no longer cause unexpected termination.

All users of libbonobo are advised to upgrade to these updated packages which fix this bug.

5.142. LIBBURN

5.142.1. RHBA-2012:1273 – libburn bug fix update

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

problem description Libburn is an open-source library for reading, mastering and writing optical discs. For now this means only CD-R and CD-RW.

Bug Fix

BZ#822906

Prior to this update, libburn library contained the "burn_write_close_track" command, which was redundant and not fully supported by all burning drives. As a consequence, the burning process CD-R or CD-RW could log errors while closing a track after the burning process, even if the data was written correctly. This update removes this redundant call.

All users of gfs-kmod are advised to upgrade to these updated packages, which fix this bug.

5.143. LIBCGROUP

5.143.1. RHBA-2012:0867 – libcgroup bug fix update

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

The libcgroup packages provides tools and libraries to control and monitor control groups.

Bug Fix

BZ#758493

Prior to this update, a bug in the libcgroup package caused admin_id and admin_gid to not be displayed correctly because cgroup control files were not differentiated from the 'tasks' file. This update fixes this problem by adding a check in the 'cgroup_fill_cgc()' function to see if a file is a 'tasks' file or not.

All users are advised to upgrade to these updated libcgroup packages, which fix this bug.

5.144. LIBDVDREAD

5.144.1. RHBA-2012:1247 – libdvdread bug fix update

Updated libdvdread packages that fix one bug is now available for Red Hat Enterprise Linux 6.

The libdvdread packages contain a simple foundation to read DVD video disks. This provides the functionality that is required to access many DVDs.

Bug Fix

BZ#842016

Prior to this update, the dvd_stat_t structure was not public. As a consequence, source code that required such structures could not be compiled. This update makes the dvd_stat_t structure public, to allow compiling code with of this type.

All users of libdvdread are advised to upgrade to these updated packages, which fix this bug.

5.145. LIBERATION-FONTS

5.145.1. RHBA-2012:0384 – liberation-fonts bug fix update

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

The liberation-fonts packages provide fonts intended to replace the three most commonly used fonts on Microsoft systems: Times New Roman, Arial, and Courier New.

Bug Fix

BZ#772165

Previously, the "fonts.dir" file provided with the liberation-fonts packages was empty. As a consequence, legacy applications were not able to make use of liberation-fonts even when the package was installed. This was because the "mkfontscale" command was run after the "mkfontdir" command. The order of running the commands has been changed and legacy applications can use liberation-fonts as expected.

All users of liberation-fonts are advised to upgrade to these updated packages, which fix this bug.

5.146. LIBEVENT

5.146.1. RHBA-2012:0968 – libevent bug fix update

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

The libevent API provides a mechanism to execute a callback function when a specific event occurs on a file descriptor or after a timeout has been reached.

Bug Fix

BZ#658051

Prior to this update, several multilib files in the optional repositories could cause conflicts. As a consequence, these files could not use both a primary and a secondary architecture. This update modifies the underlying source code so that no more multilib file conflicts occur.

All users of libevent or applications that require the libevent library are advised to upgrade to these updated packages, which fix this bug.

5.147. LIBEXIF

5.147.1. RHSA-2012:1255 – Moderate: libexif security update

Updated libexif packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

The libexif packages provide an Exchangeable image file format (Exif) library. Exif allows metadata to be added to and read from certain types of image files.

Security Fix

CVE-2012-2812, CVE-2012-2813, CVE-2012-2814, CVE-2012-2836, CVE-2012-2837, CVE-2012-2840, CVE-2012-2841

Multiple flaws were found in the way libexif processed Exif tags. An attacker could create a specially-crafted image file that, when opened in an application linked against libexif, could cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

Red Hat would like to thank Dan Fandrich for reporting these issues. Upstream acknowledges Mateusz Jurczyk of the Google Security Team as the original reporter of CVE-2012-2812, CVE-2012-2813, and CVE-2012-2814; and Yunho Kim as the original reporter of CVE-2012-2836 and CVE-2012-2837.

Users of libexif are advised to upgrade to these updated packages, which contain backported patches to resolve these issues. All running applications linked against libexif must be restarted for the update to take effect.

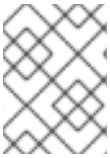
5.148. LIBGUESTFS

5.148.1. RHSA-2012:0774 – libguestfs security, bug fix and enhancement update

Updated libguestfs packages that fix one security issue, multiple 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.

The libguestfs package contains a library for accessing and modifying guest disk images.



NOTE

The libguestfs package has been upgraded to upstream version 1.16, which provides a number of bug fixes and enhancements over the previous version. (BZ#719879)

Security Fix

CVE-2012-2690

It was found that editing files with virt-edit left said files in a world-readable state (and did not preserve the file owner or Security-Enhanced Linux context). If an administrator on the host used virt-edit to edit a file inside a guest, the file would be left with world-readable permissions. This could lead to unprivileged guest users accessing files they would otherwise be unable to.

Bug Fixes

BZ#647174

When cloning, the **virt-clone** tool incorrectly adopted some of the properties of the original virtual machine image, for example, the **udev** rules for network interface: the clone was then created with a NIC identical to the NIC of the original virtual machine NIC. With this update, the **virt-sysprep** and **virt-sparsify** tools have been added to solve this problem. The **virt-sysprep** tool can erase the state from guests, and **virt-sparsify** can make guest images sparse. Users are advised to use **virt-sysprep** and **virt-sparsify** either as a replacement for or in conjunction with **virt-clone**.

BZ#789960

The **libguestfs** daemon terminated unexpectedly when it attempted to mount a non-existent disk. This happened because **libguestfs** returned an unexpected error to any program that accidentally tried to mount a non-existent disk and all further operations intended to handle such a situation failed. With this update, **libguestfs** returns an appropriate error message and remains stable in the scenario described.

BZ#790958

If two threads in one program called the **guestfs_launch()** function at the same time, an unexpected error could be returned. The respective code in the libguestfs library has been modified to be thread-safe in this scenario and the library can be used from multi-threaded programs with more than one libguestfs handle.

BZ#769359

After a block device was closed, the **udev** device manager triggered a process which re-opened the block device. Consequently, libguestfs operations occasionally failed as they rely on the disk being immediately free for the kernel to re-read the partition table. This commonly occurred with the **virt-resize** feature. With this update, the operations now wait for the **udev** action to finish and no longer fail in the scenario described.

BZ#809401

In Fedora 17, the **/bin** directory is a symbolic link, while it was a directory in previous releases. Due to this change, libguestfs could not inspect a guest with Fedora 17 and newer. With this update, the libguestfs inspection has been changed so that it now recognizes such guests as expected.

BZ#729076

Previously, libguestfs considered any disk that contained **autoexec.bat** or **boot.ini** or **ntldr** file in its root a candidate for a Windows root disk. If a guest had an HP recovery partition, libguestfs could not recognize the HP recovery partition and handled the system as being dual-boot. Consequently, some virt tools did not work as they do not support multi-boot guests. With this update, libguestfs investigates a potential Windows root disk properly and no longer recognizes the special HP recovery partition as a Windows root disk.

BZ#811673

If launching of certain appliances failed, libguestfs did not set the error string. As Python programs handling the bindings assumed that the error string was not **NULL**, the binding process terminated unexpectedly with a segmentation fault when the **g.launch()** function was called under some circumstances. With this update, the error string is now set properly on all failure paths in the described scenario and Python programs no longer terminate with a segmentation fault when calling the **g.launch()** function under these circumstances.

BZ#812092

The qemu emulator cannot open disk image files that contain the colon character (**:**). Previously, libguestfs resolved the link to the disk image before sending it to qemu. If the resolved link contained the colon character, qemu failed to run. Also, libguestfs sometimes failed to open a disk image file under these circumstances due to incorrect handling of special characters. With this update, libguestfs no longer resolves a link to a disk image before sending it to qemu and is able to handle any filenames, except for filenames that contain a colon character. Also, libguestfs now returns correct diagnostic messages when presented with a filename that contains a colon character.

Enhancements

BZ#741183

The **libguestfs** application now provides the **virt-alignment-scan** tool and updated **virt-resize**, which can diagnose unaligned partitions on a guest, so that you can fix the problem and improve the partitions' performance. For more information, refer to the virt-alignment-scan(1) and virt-resize(1) manual pages.

BZ#760221

Previously, libguestfs operations could not handle paths to HP Smart Array (cciss) devices. When the **virt-p2v** tool converted a physical machine that uses Linux software RAID devices to run in a VM, the libguestfs inspection failed to handle the paths in the /etc/fstab file. With this update, support for such cciss paths has been added and the **virt-p2v** tool is now able to successfully convert these guests.

BZ#760223

When the **virt-p2v** tool converted a physical machine that uses Linux software RAID devices to run in a VM, the libguestfs inspection failed to handle the paths in the **/etc/fstab** file. With this update, support for such RAID paths has been added and the **virt-p2v** tool is now able to successfully convert these guests.

Users of libguestfs should upgrade to these updated packages, which fix these issues and add these enhancements.

5.149. LIBGWEATHER

5.149.1. RHBA-2012:0381 – libgweather bug fix update

An updated libgweather package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The libgweather library is used by applications to access weather information from online services for numerous locations.

Bug Fix

BZ#704105

Previously, the libgweather library did not contain information for the city of Jerusalem, making it impossible to search for the city or select the city from the list of cities offered by the weather panel applet. This update adds the missing information to the libgweather database, so that users can now select Jerusalem as a valid option.

All users of libgweather are advised to upgrade to this updated package, which fixes this bug.

5.150. LIBHBAAPI

5.150.1. RHBA-2012:0847 – libhbaapi bug fix update

Updated libhbaapi packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

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 libhbaapi build environment has been upgraded to upstream version 2.2.5, which provides a number of bug fixes over the previous version. (BZ#788504)

Bug Fix

BZ#806731

Prior to this update, the hba.conf file was not marked for exclusion from verification in the libhbaapi specification file. As a consequence, the file verify function "rpm -V libhbaapi" reported an error in the hba.conf file if the file was changed. This update marks hba.conf in the spec file as "%verify(not md5 size mtime)". Now, the hba.conf file is no longer incorrectly verified.

All users of libhbaapi are advised to upgrade to these updated packages, which fix these bugs.

5.151. LIBHBALINUX

5.151.1. RHEA-2012:0848 – libhbalinux enhancement update

An updated libhbalinux package that fixes multiple bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The libhbalinux package contains 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 libhbalinux packages have been upgraded to upstream version 1.0.13, which provides a number of bug fixes and enhancements over the previous version. (BZ#719584)

All users of libhbalinux are advised to upgrade to this updated libhbalinux package, which fixes these bugs and adds these enhancements.

5.152. LIBIBVERBS-ROCEE AND LIBMLX4-ROCEE

5.152.1. RHEA-2012:0977 – libibverbs-rocee and libmlx4-rocee bug fix update

Updated libibverbs-rocee and libmlx4-rocee packages that fix one bug are now available for Red Hat Enterprise Linux High Performance Network.

The libibverbs-rocee packages provide a library to enable userspace processes to use RDMA over Converged Ethernet (RoCE) "verbs" acording to the InfiniBand Architecture Specification and the RoCE Protocol Verbs Specification. The libmlx4-rocee packages provide a device-specific driver for Mellanox ConnectX InfiniBand host channel adapters (HCAs) for the libibverbs-rocee library.

Bug Fix

BZ#805717

Prior to this update, running modprobe could, under certain circumstances, cause an infinite loop. As a consequence, the system ran out of processes and required a restart to recover. This update modifies the underlying code so that an incorrect configuration of options in the /etc/modprobe.d/libmlx4.conf file no longer requires to restart the system.



NOTE

This package is a dependency of the RDMA package and both packages have to be updated together.

All users who require RDMA are advised to upgrade to these updated packages which fix this bug.

5.153. LIBPROXY

5.153.1. RHSA-2012:1461 – Moderate: libproxy security update

Updated libproxy packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

libproxy is a library that handles all the details of proxy configuration.

Security Fix

CVE-2012-4505

A buffer overflow flaw was found in the way libproxy handled the downloading of proxy auto-configuration (PAC) files. A malicious server hosting a PAC file or a man-in-the-middle attacker could use this flaw to cause an application using libproxy to crash or, possibly, execute arbitrary code, if the proxy settings obtained by libproxy (from the environment or the desktop environment settings) instructed the use of a PAC proxy configuration.

This issue was discovered by the Red Hat Security Response Team.

Users of libproxy should upgrade to these updated packages, which contain a backported patch to correct this issue. All applications using libproxy must be restarted for this update to take effect.

5.154. LIBREOFFICE

5.154.1. RHSA-2012:1135 – Important: libreoffice security update

Updated libreoffice packages that fix multiple security issues are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

LibreOffice is an open source, community-developed office productivity suite. It includes the key desktop applications, such as a word processor, spreadsheet application, presentation manager, formula editor, and a drawing program.

Security Fix

CVE-2012-2665

Multiple heap-based buffer overflow flaws were found in the way LibreOffice processed encryption

information in the manifest files of OpenDocument Format files. An attacker could provide a specially-crafted OpenDocument Format file that, when opened in a LibreOffice application, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

Upstream acknowledges Timo Warns as the original reporter of these issues.

All LibreOffice users are advised to upgrade to these updated packages, which contain backported patches to correct these issues. All running instances of LibreOffice applications must be restarted for this update to take effect.

5.155. LIBSELINUX

5.155.1. RHBA-2012:0907 – libselinux bug fix update

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

The libselinux packages contain the core library of an SELinux system. The libselinux library provides an API for SELinux applications to get and set process and file security contexts, and to obtain security policy decisions. It is required for any applications that use the SELinux API, and used by all applications that are SELinux-aware.

Bug Fix

BZ#717147

While the libselinux library was waiting on a netlink socket, if the socket received an EINTR signal, it returned an error which could cause applications like dbus to fail. With this update, the library now retries the netlink socket when it receives an EINTR signal, rather than failing.

All users of libselinux are advised to upgrade to these updated packages, which fix this bug.

5.155.2. RHEA-2013:0808 – libselinux enhancement update

Updated libselinux packages that add one enhancement are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The libselinux packages contain the core library of an SELinux system. The libselinux library provides an API for SELinux applications to get and set process and file security contexts, and to obtain security policy decisions. It is required for any applications that use the SELinux API, and used by all applications that are SELinux-aware.

Enhancement

BZ#956982

Previously, a substitution of the "/" directory was not directly possible. With this update, support for a substitution of the root directory has been added to allow proper labeling of all directories and files under an alternative root directory.

Users of libselinux are advised to upgrade to these updated packages, which adds this enhancement.

5.156. LIBSERVICELOG

5.156.1. RHBA-2012:0988 – libservicelog bug fix update

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

The libservicelog packages provide a library for logging service-related events to the servicelog database, and a number of command-line utilities for viewing the contents of the database.

Bug Fix

BZ#814171

Prior to this update, the "servicelog_manage --truncate" command cleared only the "events" table. As a consequence, the other tables were not deleted. This update modifies the option "servicelog_event_delete()" function so that all rows in all tables associated with a deleted event are correctly deleted. Now, the tables for "callouts", "os", "rtas", "enclosure" and "repair_actions" are cleared together with the "events" table.

All users of libservicelog are advised to upgrade to these updated packages, which fix this bug.

5.157. LIBSSH2

5.157.1. RHBA-2012:1048 – libssh2 bug fix update

Updated libssh2 packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The libssh2 packages provide a library that implements the SSH2 protocol.

Bug Fix

BZ#834211

Previously, libssh2 incorrectly returned the LIBSSH2_ERROR_EAGAIN error code when operating in blocking mode. The error code is used by libssh2 internally to initiate a blocking operation on a socket. The error code was, however, not properly cleared on success and leaked through the public API of libssh2. An upstream patch has been applied to clear the error code prior to initiating the blocking operation, and libssh2 no longer returns LIBSSH2_ERROR_EAGAIN when operating in blocking mode.

All users of libssh2 are advised to upgrade to these updated packages, which fix this bug. After installing these updated packages, all running applications using libssh2 have to be restarted for this update to take effect.

5.158. LIBTAR

5.158.1. RHBA-2012:0462 – libtar bug fix update

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

The libtar package contains a C library for manipulating tar archives. The library supports both the strict POSIX tar format and many of the commonly used GNU extensions.

Bug Fix

BZ#729009

Previously, the build system configuration files included in the libtar package were incompatible with the way the rpmbuild tool extracts debugging information from the binaries installed to the rpm build root during the build of a package. As a consequence, the libtar-debuginfo package did not contain debugging information. A patch has been applied to address this issue, and the libtar-debuginfo package now contains the appropriate content.

All users of libtar are advised to upgrade to these updated packages, which fix this bug.

5.159. LIBTIFF

5.159.1. RHSA-2012:1590 – Moderate: libtiff security update

Updated libtiff packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.

Security Fixes

CVE-2012-4447

A heap-based buffer overflow flaw was found in the way libtiff processed certain TIFF images using the Pixar Log Format encoding. An attacker could create a specially-crafted TIFF file that, when opened, could cause an application using libtiff to crash or, possibly, execute arbitrary code with the privileges of the user running the application.

CVE-2012-5581

A stack-based buffer overflow flaw was found in the way libtiff handled DOTRANGE tags. An attacker could use this flaw to create a specially-crafted TIFF file that, when opened, would cause an application linked against libtiff to crash or, possibly, execute arbitrary code.

CVE-2012-3401

A heap-based buffer overflow flaw was found in the tiff2pdf tool. An attacker could use this flaw to create a specially-crafted TIFF file that would cause tiff2pdf to crash or, possibly, execute arbitrary code.

CVE-2012-4564

A missing return value check flaw, leading to a heap-based buffer overflow, was found in the ppm2tiff tool. An attacker could use this flaw to create a specially-crafted PPM (Portable Pixel Map) file that would cause ppm2tiff to crash or, possibly, execute arbitrary code.

The CVE-2012-5581, CVE-2012-3401, and CVE-2012-4564 issues were discovered by Huzaifa Sidhpurwala of the Red Hat Security Response Team.

All libtiff users should upgrade to these updated packages, which contain backported patches to resolve these issues. All running applications linked against libtiff must be restarted for this update to take effect.

5.159.2. RHSA-2012:1054 – Important: libtiff security update

Updated libtiff packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.

Security Fixes

CVE-2012-2088

libtiff did not properly convert between signed and unsigned integer values, leading to a buffer overflow. An attacker could use this flaw to create a specially-crafted TIFF file that, when opened, would cause an application linked against libtiff to crash or, possibly, execute arbitrary code.

CVE-2012-2113

Multiple integer overflow flaws, leading to heap-based buffer overflows, were found in the tiff2pdf tool. An attacker could use these flaws to create a specially-crafted TIFF file that would cause tiff2pdf to crash or, possibly, execute arbitrary code.

All libtiff users should upgrade to these updated packages, which contain backported patches to resolve these issues. All running applications linked against libtiff must be restarted for this update to take effect.

5.160. LIBUNISTRING

5.160.1. RHBA-2012:0887 – libunistring bug fix update

An updated libunistring package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The libunistring package contains a portable C library that implements the UTF-8, UTF-16 and UTF-32 Unicode string types, together with functions for character processing (names, classifications, and properties) and functions for string processing (iteration, formatted output, width, word breaks, line breaks, normalization, case folding, and regular expressions).

Bug Fix

BZ#732017

Previously, when calling the malloc() function, no check for the returned pointer was performed to find out whether memory was successfully allocated. Therefore, if a null pointer was returned, this could cause the libunistring library to misbehave in low-memory situations. This update adds the missing check to properly handle such situations.

All users of libunistring are advised to upgrade to this updated package, which fixes this bug.

5.161. LIBUSB1

5.161.1. RHBA-2012:0759 – libusb1 bug fix and enhancement update

Updated libusb1 package that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The libusb1 packages provide a way for applications to access USB devices.

The libusb1 packages have been upgraded to upstream version 1.0.9, which provides a number of bug fixes and enhancements over the previous version. In addition, this update adds a new API needed for support of the SPICE (The Simple Protocol for Independent Computing Environments) USB redirection. (BZ#758094)

All users of libusb1 are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.162. LIBUSER

5.162.1. RHBA-2012:0455 – libuser bug fix update

Updated libuser packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The libuser library implements a standardized interface for manipulating and administering user and group accounts. The library uses pluggable back-ends to interface to its data sources. Sample applications modeled after those included with the shadow password suite are included.

Bug Fixes

BZ#670151

When creating a user account in Lightweight Directory Access Protocol (LDAP), the libuser library used the value of the "gecos" attribute as the default value of the "cn" attribute. When the "gecos" attribute was empty, this made the value of "cn" invalid, and the creation of the user account failed. With this update, the user name of the account is stored in the "cn" attribute if the "gecos" attribute is empty, thus allowing successful creation of the user account.

BZ#724987

When populating a home directory by copying files from the /etc/skel directory, libuser ignored the "set user-id" and "set group-id" flags. This made it impossible to set up group-shared directories in a home directory. With this update, the "set user-id" and "set group-id" flags are preserved.

BZ#788521

Previously, when searching for the user or group account information in files of certain sizes, the libuser library could terminate unexpectedly with a segmentation fault. A patch has been applied to address this issue, and crashes no longer occur in the aforementioned scenario.

All users of libuser are advised to upgrade to these updated packages, which fix these bugs.

5.163. LIBVIRT-CIM

5.163.1. RHBA-2012:0757 – libvirt-cim bug fix and enhancement update

An updated libvirt-cim package that fixes various bugs and adds multiple enhancements is now available for Red Hat Enterprise Linux 6.

The libvirt-cim package contains 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.

The libvirt-cim package has been upgraded to upstream version 0.6.1, which provides a number of bug fixes and enhancements over the previous version. (BZ#739154)

Bug Fix

BZ#799037

Previously, the libvirt-cim package required as its dependency the tog-pegasus package, which contains the OpenPegasus Web-Based Enterprise Management (WBEM) services. This is, however, incorrect as libvirt-cim should not require specifically tog-pegasus but any CIM server. With this update, libvirt-cim has been changed to require cim-server instead. The spec files of libvirt-cim and sblim-sfcb have been modified appropriately and libvirt-cim now uses either of the packages as its dependency.

Enhancements

BZ#633338

Extension for Quality-of-Service (QoS) networking has been added.

BZ#739153

Support for domain events has been added.

BZ#739156

Extensions for networking of Central Processing Unit (CPU) shares have been added.

All libvirt-cim users are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.164. LIBVIRT-JAVA

5.164.1. RHBA-2012:1075 – libvirt-java bug fix

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

The libvirt-java packages provide Java bindings to use libvirt, which is the virtualization API to manage and interact with virtualization capabilities.

Bug Fix

BZ#836920

Prior to this update, the "setSchedulerParameters()" method for domains did not work as expected because the jna conversion was failing. This update modifies the conversion process. Now, the given parameter is used as expected.

All users of libvirt-java are advised to upgrade to these updated packages, which fix this bug.

5.165. LIBVIRT-QMF

5.165.1. RHBA-2012:1001 – libvirt-qmf bug fix update

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

The libvirt-qmf packages provide an interface with libvirt using Qpid Management Framework (QMF), which utilizes the Advanced Message Queuing Protocol (AMQP). AMQP is an open standard application layer protocol providing reliable transport of messages.

Bug Fix

BZ#830087

The libvirt-qmf packages included in Red Hat Enterprise Linux 6.2 were of version 0.3.0-7, whereas the packages in Red Hat Enterprise Linux 6.3 were of version 0.3.0-6. To prevent possible problems with upgrading, Red Hat Enterprise Linux 6.3 now uses libvirt-qmf version 0.3.0-8.

All users of libvirt-qmf are advised to upgrade to these updated packages, which fix this bug.

5.165.2. RHBA-2012:0983 – libvirt-qmf bug fix update

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

The libvirt-qmf packages provide an interface with libvirt using Qpid Management Framework (QMF), which utilizes the Advanced Message Queuing Protocol (AMQP). AMQP is an open standard application layer protocol providing reliable transport of messages.

Bug Fix

BZ#806950

Prior to this update, Qpid APIs used the libpidclient and libpidcommon libraries, which were not application binary interface (ABI) stable. This update removes these dependencies so that Qpid rebuilds do not affect the libvirt-qmf packages.

All users of libvirt-qmf are advised to upgrade to these updated packages, which fix this bug.

5.166. LIBVIRT

5.166.1. RHBA-2012:1595 – libvirt bug fix update

Updated libvirt packages that fix several bugs 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#877024

The AMD FX series processors contain "modules" which are reported by the kernel as both threads and cores. Previously, the processor topology detection code in libvirt was unable to detect these modules. Consequently, libvirt reported twice the actual number of processors. With this update,

topologies that add up to the total number of processors reported by the system are properly reported even though the actual topology has to be checked in the output of the virCapabilities() function.

Note that the capability output for topology detection purposes should be used due to performance reasons. The NUMA topology has high impact on performance but the impact of the physical topology can differ from that.

BZ#884713

Whenever the virDomainGetXMLDesc() function was executed on a domain that was unresponsive, the call also became unresponsive. With this update, QEMU sends the BALLOON_CHANGE event when memory usage on a domain changes so that virDomainGetXMLDesc() no longer has to query an unresponsive domain. As a result, virDomainGetXMLDesc() calls no longer hang in the described scenario.

All users of libvirt are advised to upgrade to these updated packages, which fix these bugs.

5.166.2. RHSA-2013:0199 – Important: libvirt security update

Updated libvirt packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

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.

Security Fix

CVE-2013-0170

A flaw was found in the way libvirtd handled connection cleanup (when a connection was being closed) under certain error conditions. A remote attacker able to establish a read-only connection to libvirtd could use this flaw to crash libvirtd or, potentially, execute arbitrary code with the privileges of the root user.

This issue was discovered by Tingting Zheng of Red Hat.

All users of libvirt are advised to upgrade to these updated packages, which contain a backported patch to correct this issue. After installing the updated packages, libvirtd will be restarted automatically.

5.166.3. RHSA-2012:1359 – Moderate: libvirt security and bug fix update

Updated libvirt packages that fix one security issue and multiple 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

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.

Security Fix

CVE-2012-4423

A flaw was found in libvirtd's RPC call handling. An attacker able to establish a read-only connection to libvirtd could use this flaw to crash libvirtd by sending an RPC message that has an event as the RPC number, or an RPC number that falls into a gap in the RPC dispatch table.

This issue was discovered by Wenlong Huang of the Red Hat Virtualization QE Team.

Bug Fixes

BZ#858988

When the host_uuid option was present in the libvirtd.conf file, the augeas libvirt lens was unable to parse the file. This bug has been fixed and the augeas libvirt lens now parses libvirtd.conf as expected in the described scenario.

BZ#859376

Disk hot plug is a two-part action: the qemuMonitorAddDrive() call is followed by the qemuMonitorAddDevice() call. When the first part succeeded but the second one failed, libvirt failed to roll back the first part and the device remained in use even though the disk hot plug failed. With this update, the rollback for the drive addition is properly performed in the described scenario and disk hot plug now works as expected.

BZ#860720

When a virtual machine was started with an image chain using block devices and a block rebase operation was issued, the operation failed on completion in the blockJobAbort() function. This update relabels and configures cgroups for the backing files and the rebase operation now succeeds.

All users of libvirt are advised to upgrade to these updated packages, which contain backported patches to correct these issues. After installing the updated packages, libvirtd will be restarted automatically.

5.166.4. RHSA-2012:1202 – Moderate: libvirt security and bug fix update

Updated libvirt 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

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.

Security Fix

CVE-2012-3445

A flaw was found in libvirtd's RPC call handling. An attacker able to establish a read-only connection to libvirtd could trigger this flaw with a specially-crafted RPC command that has the number of parameters set to 0, causing libvirtd to access invalid memory and crash.

Bug Fixes

BZ#847946

Previously, repeatedly migrating a guest between two machines while using the tunnelled migration could cause the libvirt daemon to lock up unexpectedly. The bug in the code for locking remote drivers has been fixed and repeated tunnelled migrations of domains now work as expected.

BZ#847959

Previously, when certain system locales were used by the system, libvirt could issue incorrect commands to the hypervisor. This bug has been fixed and the libvirt library and daemon are no longer affected by the choice of the user locale.

All users of libvirt are advised to upgrade to these updated packages, which contain backported patches to correct these issues. After installing the updated packages, libvirtd will be restarted automatically.

5.166.5. RHBA-2012:1484 – libvirt bug fix update

Updated libvirt packages that fix several bugs 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#868972

When libvirt could not find a suitable CPU model for a host CPU, it would not provide the CPU topology in host capabilities even though the topology was detected correctly. Consequently, applications that work with the host CPU topology but not with the CPU model could not see the topology in host capabilities. With this update, the host capabilities XML description contains the host CPU topology even if the host CPU model is unknown.

BZ#869650

Previously, the fixed limit for the maximum size of an RPC message that could be supplied to the libvirtd daemon (65536 bytes) was not always sufficient. Consequently, messages that were longer than that could be dropped, leaving a client unable to fetch important data. With this update, the buffer for incoming messages has been made dynamic and libvirtd now allocates as much memory as is needed for a given message, thus allowing to send much bigger messages.

BZ#869723

Prior to this update, libvirt used an unsuitable detection procedure to detect NUMA and processor topology of a system. Consequently, topology of some advanced multi-processor systems was detected incorrectly and management applications could not utilize the full potential of the system. Now, the detection has been improved and the topology is properly recognized even on modern systems.

BZ#873292

Under certain circumstances, the iohelper process failed to write data to disk while saving a domain and kernel did not report an out-of-space error (ENOSPC). With this update, libvirt calls the fdatasync() function in the described scenario to force the data to be written to disk or catch a write error. As a result, if a write error occurs, it is now properly caught and reported.

BZ#874235

Certain operations in libvirt can be done only when a domain is paused to prevent data corruption. However, if a resuming operation failed, the management application was not notified since no event was sent. This update introduces the VIR_DOMAIN_EVENT_SUSPENDED_API_ERROR event and management applications can now keep closer track of domain states and act accordingly.

BZ#875770

Libvirt allows users to cancel an ongoing migration. Previously, if an attempt to cancel the migration was made in the migration preparation phase, qemu missed the request and the migration was not canceled. With this update, the virDomainAbortJob() function sets a flag when a cancel request is made and this flag is checked before the main phase of the migration starts. As a result, a migration can now be properly canceled even in the preparation phase.

BZ#875788

When a qemu process is being destroyed by libvirt, a clean-up operation frees some internal structures and locks. However, since users can destroy qemu processes at the same time, libvirt holds the qemu driver lock to protect the list of domains and their states, among other things. Previously, a function tried to set up the qemu driver lock when it was already up, creating a deadlock. The code has been modified to always check if the lock is free before attempting to set it up, thus fixing this bug.

All users of libvirt are advised to upgrade to these updated packages, which fix these bugs.

5.166.6. RHBA-2012:1095 – libvirt bug fix update

Updated libvirt packages that fix two bugs 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#836916

Previously, repeatedly attaching and detaching a PCI device to a guest domain could cause the libvirt daemon to terminate unexpectedly. The erroneous structure free operation at the root of this bug has been fixed and repeated attach and detach actions of a PCI device now work as expected.

BZ#836919

On certain NUMA architectures, libvirt was failing to process and expose the NUMA topology, possibly leading to performance degradation. These updated packages now correctly parse and expose the NUMA topology on such machines and make the correct CPU placement, thus avoiding the performance degradation.

BZ#838819

When using the sanlock daemon for locking resources used by a domain, if such a resource was read-only, the locking attempt failed. Consequently, it was impossible to start a domain with a CD-ROM drive. This bug has been fixed and sanlock can now be properly used with read-only devices.

All users of libvirt are advised to upgrade to these updated packages, which fix these bugs.

5.166.7. RHBA-2012:1000 – libvirt bug fix update

Updated libvirt packages that fix two bugs 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#827050

Closing a file descriptor multiple times could, under certain circumstances, lead to a failure to execute the qemu-kvm binary. As a consequence, a guest failed to start. A patch has been applied to address this issue, so that the guest now starts successfully.

BZ#832184

Libvirt 0.9.10 has added support for keepalive checking to detect broken connections between the client and the server. However, due to bugs in the implementation this could have caused a failure of service and disconnection, for example, during parallel migrations. The keepalive support is now disabled by default and random disconnections no longer occur.

All users of libvirt are advised to upgrade to these updated packages, which fix these bugs.

5.166.8. RHSA-2012:0748 – libvirt security, bug fix, and enhancement update

Updated libvirt packages that fix one security issue, multiple 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.

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.



NOTE

The libvirt packages have been upgraded to upstream version 0.9.10, which provides a number of bug fixes and enhancements over the previous version. (BZ#752433)

Security Fix

CVE-2012-2693

Bus and device IDs were ignored when attempting to attach multiple USB devices with identical vendor or product IDs to a guest. This could result in the wrong device being attached to a guest, giving that guest root access to the device.

Bug Fixes

BZ#754621

Previously, libvirt incorrectly released resources in the macvtap network driver in the underlying code for QEMU. As a consequence, after an attempt to create a virtual machine failed, a macvtap device that was created for the machine could not be deleted from the system. Any virtual machine using the same MAC address could not be created in such a case. With this update, an incorrect function call has been removed, and macvtap devices are properly removed from the system in the scenario described.

BZ#742087

Under certain circumstances, a race condition between asynchronous jobs and query jobs could occur in the QEMU monitor. Consequently, after the QEMU guest was stopped, it failed to start again with the following error message:

```
error: Failed to start domain [domain name]
error: Timed out during operation cannot acquire state change lock
```

With this update, libvirt handles this situation properly, and guests now start as expected.

BZ#769500

Previously, libvirt defined a hard limit for the maximum number of virtual machines (500) in Python bindings. As a consequence, the vdsmd daemon was unable to properly discover all virtual machines on a system with more than 500 guests. With this update, the number of virtual machines is now determined dynamically and vdsmd correctly discovers all virtual machines.

BZ#739075

Previously, it was not possible to cancel all migration-family commands (for example, it was possible to cancel the "virsh migration" command, but not the "virsh dump" command). This update implements a mechanism used for "virsh migration" also for the other commands, so it is now possible to cancel these commands.

BZ#773667

Previously, libvirt was unable to verify if there were multiple active PCI devices on the same I/O bus. As a consequence, the "virsh attach-device" command failed even if such a device had already been detached from the host. With this update, libvirt properly checks for active devices on the same PCI I/O bus. Users can now attach devices to a guest successfully if the devices on the same bus are detached from the host.

BZ#701654

When the libvirt's virDomainDestroy API is shutting down the qemu process, the API first sends the SIGTERM signal, then waits for 1.6 seconds and, if the process is still running, the API sends the SIGKILL signal. Previously, this could lead to data loss because the guest running in QEMU did not have time to flush its disk cache buffers before it was unexpectedly killed. This update provides a new flag, "VIR_DOMAIN_DESTROY_GRACEFUL". If this flag is set in the call to virDomainDestroyFlags, SIGKILL is not sent to the qemu process; instead, if the timeout is reached and the qemu process still exists, virDomainDestroy returns an error. It is recommended that management applications always

first call virDomainDestroyFlags with VIR_DOMAIN_DESTROY_GRACEFUL. If that fails, then the application can decide if and when to call virDomainDestroyFlags again without VIR_DOMAIN_DESTROY_GRACEFUL.

BZ#784151

The localtime_r() function used in the libvirt code was not async-signal safe, which caused child processes to enter a deadlock when attempting to generate a log message. As a consequence, the virsh utility became unresponsive. This update applies backported patches and adds a new API for generating log time stamps in an async-signal safe manner. The virsh utility no longer hangs under these circumstances.

BZ#785269

Previously, if the libvirt package was built with Avahi support, libvirt required the avahi package to be installed on the system as a prerequisite for its own installation. If the avahi package could not be installed on the system due to security concerns, installation of libvirt failed. This update modifies the libvirt.spec file to require only the avahi-libs package. The libvirt package is now successfully installed and libvirtd starts as expected.

BZ#639599

The schema for the XML files contained stricter rules than those that were actually enforced by libvirt. As a consequence, validation tools failed to validate guest XML files that contained special characters in the guest's name even if libvirt accepted the XML file. With this update, the XML schema now allows arbitrary strings with no limitation, leaving the enforcement of rules to the hypervisor driver. As a result, users are now able to validate these XML files.

BZ#785164

Previously, the libxml2 tool did not parse IPv6 URIs as expected. As a consequence, attempting to establish an IPv6 connection through SSH failed, because an invalid IPv6 address was used. A patch has been applied to address this problem and IPv6 connections can now be established successfully in this scenario.

BZ#625362

Previously, the libvirt-guests init script executed operations on guests serially. Consequently, on machines with many guests, the shutdown process took a long time because guests were waiting for other guests to be shut down. The libvirt-guests init script was modified to enable parallel operation on domains, which reduces the time of the shutdown process of the host. Now, guests start and shut down in parallel, and utilize the host system's resources more efficiently.

BZ#783968

When migrating a QEMU virtual machine and using SPICE for a remote display, the migration was failing and the display was erratic under certain circumstances. This was happening because with the incoming migration connection open, QEMU was unable to accept any other connections on the target host. With this update, the underlying code has been modified to delay the migration connection until the SPICE client is connected to the target destination. The guest virtual machines can now be successfully migrated without disrupting the display during the migration.

BZ#701106

Previously, migration of a virtual machine failed if the machine had an ISO image attached as a CD-ROM drive and the ISO domain was inactive. With this update, libvirt introduces the new startupPolicy attribute for removable devices, which allows marking CD-ROM and diskette drives as optional. With this option, virtual machines can now be started or migrated without a removable drive if the source image is inaccessible.

BZ#725373

When a destination host lost network connectivity while a domain was being migrated to it, the migration process could not be canceled. This update implements an internal keep-alive protocol, which is able to detect broken connections or blocked libvirt daemons. When such a situation is detected during migration, libvirt now automatically cancels the process.

BZ#729694

With certain combinations of IDE and VirtIO disks, a guest operating system did not boot after the installation process. This happened because the order of disks in which they were presented to the guest during the installation was different from the order used after the installation. As a result, the system could have been installed on a disk which was not used as the primary bootable disk. With this update, libvirt makes sure that the order in which disks are presented to a guest operating system during the installation is the correct order that will be used later once the guest operating system is installed.

BZ#729940

Previously, libvirt did not provide any way to prevent multiple clients from accessing a console device. When two clients connected to a single console of a guest, the connections entered a race condition on reading data from the console device. Each of the connections only got a fragment of the data and that fragment was not copied to the other connection. This rendered the terminal unusable to all the simultaneous connections. With this update, when opening a console, a check is performed to ensure that only one client is connected to it at a given point in time. If such a session is locked, a new connection has the ability to disconnect previous console sessions. Users are now able to safely access the console and disconnect inactive sessions to take control of a guest in case the console is accidentally left connected.

BZ#769503

Virtualization hosts can have thousands of CPUs and run a thousand guests, and libvirt should be capable of controlling all of them. However, libvirt was not able to do so, and the limit was below 1000, and users were therefore unable to fully utilize their hardware. With this update, the array of file descriptors which is passed to the child process is now allocated dynamically and can handle as many file descriptors as possible. Moreover, init and startup scripts have been changed so that the maximum limit of open files can be overridden for the libvirtd daemon. Users can now fully utilize their hardware and run as many guests as they require.

BZ#746666

Due to several problems with security labeling, libvirtd became unresponsive when destroying multiple guest domains with disks on an unreachable NFS storage device. This update fixes the security labeling problems and libvirtd no longer hangs under these circumstances.

BZ#795305

When live migration of a guest was terminated abruptly (using the Ctrl+C key combination), the libvirt daemon could have failed to accept any future migration request of that guest with the following error message:

```
error: Timed out during operation: cannot acquire state change lock
```

This update adds support for registering cleanup callbacks which are called for a domain when a connection is closed. The migration API is more robust to failures, and if a migration process is terminated, it can be restarted with a subsequent command.

BZ#752255

Previously, libvirt's implementation of nwfilter attempted to execute a temporary file generated directly in the /tmp/ directory, which failed if /tmp/ was mounted with the "noexec" options for security reasons. The implementation of nwfilter has been improved to avoid the need for a temporary file altogether, so it is no longer necessary for libvirt to modify or use files in the /tmp/ directory.

BZ#575160

Prior to this update, QEMU did not provide a notify mechanism when a block device tray status was changed. As a consequence, libvirt was unable to determine if the block data medium was ejected or was not present inside a guest. If the medium was ejected inside a guest, libvirt started the guest with the media being still present when migrating, saving and restoring the guest. This update introduces a new XML attribute for removable disk devices to represent and update the tray status.

BZ#758026

Under certain circumstances, a rare race condition between the poll() event handler and the dmidecode utility could occur. This race could result in dmidecode waiting indefinitely to perform a read operation on the already closed file descriptor. As a consequence, it was impossible to perform any tasks for virtualized guests using the libvirtd management daemon, or perform certain tasks using the virt-manager utility, such as creating a new virtual machine. This update modifies the underlying code so that the race condition no longer occurs and libvirtd and virt-manager work as expected.

BZ#758870

The libvirtd daemon could become unresponsive when starting the QEMU driver because the dmidecode tool needed a lot of time to process a large amount of data. It was consequently impossible to connect to the QEMU driver. The underlying source code has been modified to properly handle the POLLHUP event, so that users can now connect to the QEMU driver successfully in this scenario.

BZ#767333

The management application can request a guest to shut down or reboot. However, this was previously implemented by issuing Advanced Configuration and Power Interface (ACPI) events to a guest which could have ignored them. Consequently, the management application was unable to reboot such a guest. This update implements support for the guest-agent application that runs on a guest and calls the "shutdown" or "reboot" command when required. This means that a guest can be shut down or rebooted even when the guest ignores the ACPI events.

BZ#754128

When shutting down, a virtual machine changed its status from the "Up" state to the "Paused" state before it was shutdown. The "Paused" state represented the state when the guest had been already stopped, but QEMU was flushing its internal buffers and was waiting for libvirt to kill it. This state change confused users so this update adds respective events and modifies libvirt to use the "shutdown" state. A virtual machine now moves from "Up" to "Powering Down" and then to the "Down" state.

BZ#733587

If a domain failed to start, the host device for the domain was re-attached to the host regardless of whether the device was used by another domain. The underlying source code has been modified so that the device that is being used by another domain is not re-attached.

BZ#726174

Differences between the Red Hat Enterprise Linux and Debian implementations of the "nc"

command, such as the presence or absence of the "-q" option, could lead to various problems. For example attempting to use a remote connection from a client expecting certain behavior to a server providing another behavior could fail on reconnection. With this update, libvirt probes capabilities of the "nc" command, and uses the appropriate options of the server even if the options differ from the "nc" on the client, which allows for successful interaction between either type of operating system.

BZ#771603

In Red Hat Enterprise Linux 6.2, libvirt unconditionally reserved PCI address 0:0:2.0 for a VGA adapter. Any domain that was created using an earlier version of libvirt with no VGA adapter and had another PCI device attached at this address could not be started. With this update, libvirt does not automatically use this PCI address for any device except for a VGA adapter. However, other devices can be attached at this address explicitly (either by the user or by using an older version of libvirt) and libvirt does not forbid to start domains with such devices. Thus, domains that could not be migrated from Red Hat Enterprise Linux 6.1 to 6.2 can be migrated from Red Hat Enterprise Linux 6.1 to 6.3.

BZ#782457

Previously, QEMU only offered the ability to perform a live snapshot of one disk at a time, but with no rollback functionality if the snapshot process failed. With this update, libvirt has been enhanced to take advantage of QEMU improvements that guarantee that either all disks have a successful snapshot, or that the failure is detected before any change which cannot be rolled back is made. This is easier for management applications performing a live disk snapshot of a guest with multiple disks.

BZ#697808

Parsing an XML file containing an incorrect root element caused an incorrect and confusing error to be displayed. The error message has been modified to display proper and detailed information about the problem when the user provides an incorrect XML file.

BZ#815206

If the umask used when starting init services was set to mask the executable or the search bit for other users, KVM virtual machines that were explicitly configured to use the "hugepages" mechanism could fail to start because the QEMU user was unable to access the directory that libvirt had created for QEMU in the hugetlbfs file system. This was because while the directory itself was owned by QEMU, its parent directory was not searchable by QEMU. To prevent this problem, when creating the parent directory, libvirt now makes sure that the parent directory is searchable by anyone regardless of umask settings.

BZ#796526

Previously, libvirt returned guest memory values in kibibytes (multiples of 1024), but with no indication of the scale. Furthermore, the libvirt documentation referred to kilobytes (multiples of 1000). Also, QEMU used mebibytes (multiples of 1024*1024) and these differences in scale could result in users making mistakes, such as giving a guest 1000 times less memory than planned, with a failure mode that was not easy to diagnose. Now the output is clear on the unit used, and the input allows users to use other units that can be more convenient.

BZ#619846

Previously, the qemu monitor command "query-migrate" did not return any error message when a problem occurred. Consequently, libvirt produced the "Migration unexpectedly failed" error message, which did not provide the proper information about the problem. The "fd:" protocol is now used to retrieve and produce the exact error message when a problem occurs.

BZ#624447

In some configurations, log messages similar to the following could be reported to libvirt or Red Hat Enterprise Virtualization users:

```
warning : virDomainDiskDefForeachPath:7654 : Ignoring open failure on xxx.xxx
```

These messages were harmless and could be safely ignored. With this update, the messages are no longer reported unless a problem occurs.

BZ#638633

Previously, libvirt and virsh ignored any script file given in the specification for a network interface of a type that did not actually use script files. To avoid confusion, this is now explicitly prohibited, an error is logged, and attempting to specify a script file for an interface type that does not support script files fails.

BZ#726771

This update provides improvements in reporting errors in XML file parsing, which makes identifying of errors easier.

BZ#746111

The libvirt package was missing a dependency on the avahi-libs package. The dependency is required due to libvirt linking in libavahi-client for mDNS support. As a consequence, the libvirtd daemon failed to start if the libvirt package was installed on the system without the avahi-libs package. With this update, the dependency on avahi-libs is now defined in the libvirt.spec file, and avahi-libs is installed along with libvirt.

BZ#802856

In previous versions of Red Hat Enterprise Linux, a "hostdev" device could be hot plugged to a guest, but making that device persistent across restarts of the guest required separately editing the guest configuration. This update adds support for persistent hot plug of "hostdev" devices, both to the libvirt API and to the virsh utility.

BZ#806633

Previously, attempting to migrate a server from a bridge network to direct network could fail when using libvirt with a virtio network interface. With this update, if a virtual guest created using the tools in Red Hat Enterprise Linux 6.2 or earlier is started on a host running Red Hat Enterprise Linux 6.3 with the vhost-net driver module loaded, and if that guest has a virtio network interface that uses macvtap, the "merge receive buffers" feature of the virtio driver is disabled. Compatibility with Red Hat Enterprise Linux 6.2 hosts is preserved and migration no longer fails under these circumstances.

Enhancements

BZ#761005

With this update, libvirt now supports for the latest Intel processors and new features these processors include.

BZ#767364

With this update, libvirt now supports family 15h microarchitecture AMD processors.

BZ#643373

Now, libvirt is capable of controlling the state (up or down) of a link of the guest virtual network

interfaces. This allows users to perform testing and simulation as though plugging and unplugging the network cable from the interface. This feature also lets users isolate guests in case any issues arise.

BZ#691539

This update adds the ability to assign an SR-IOV (Single Root I/O Virtualization) network device Virtual Functions (VF) to a guest using the "interface" element rather than the "hostdev" element. This gives the user the opportunity to specify a known or fixed MAC address (**<mac address='xx:xx:xx:xx:xx:xx'/>**).

BZ#638506

Previously, the only way to perform storage migration was to stop a guest, edit the XML configuration file, and restart the guest. This led to a downtime that could have lasted several minutes. With this update, it is now possible to perform live storage migration with minimal guest downtime. This is ensured by new libvirt API flags to the virDomainStorageRebase() function, which map to new QEMU features.

BZ#693842

Previously, libvirt was able to notify a switch capable of the 802.1Qbg standard about changes in the guest network interface configuration, but there was no way for the switch to notify libvirt. This update provides extended support for libvirt synchronization with the lldpad daemon. As a result, if there are changes in the network infrastructure that require libvirt to re-associate the guest's interface, libvirt is informed and can take the proper action.

BZ#782034

With this update, libvirt supports a new model for the Small Computer System Interface (SCSI) controller, virtio-scsi.

BZ#713170

With this update, "fabric_name" of the "fc_host" class is exposed, so that users can see which fabric the virtual host bus adapter (vHBA) is connected to.

BZ#715019

This update introduces a new API, which allows the management system to query the disk latency using libvirt.

BZ#725013

It is sometimes required not only to delete a domain's storage but also overwrite the data to make sure sensitive data are no longer readable. This update introduces a new API, that allows users to erase the storage and use various wiping patterns.

BZ#769930

With this update, libvirt supports dynamic NUMA tuning, so that significant processes can be pre-bound to nodes with sufficient available resources.

BZ#740375

Previously, when doing disk snapshots, the guest had to be paused in order to avoid writing of data. Otherwise, the data could be corrupted. A new utility, guest-agent, has been introduced, and allows to freeze disks or file systems from inside the guest. It is no longer needed to pause the guest. However, disk write operations are delayed until the snapshot is completed.

BZ#768450

Previously, no mappings were specified for the "cpu64-rhel*" CPU models found in QEMU and therefore they could not be used. These mappings have been added with this update.

BZ#754073

Previously, it was not possible to see the memory used by the qemu-kvm process using only the virsh utility. The API call that reports the domain memory statistics has been modified to show this value. The value is now displayed when running the "virsh dommemstat" command.

BZ#533138

This update adds support for hot plugging and unplugging processors. It is now possible to add CPUs to guests and remove them as needed, without shutting down the guest.

BZ#713932

This update introduces a new virsh command, "change-media", which makes it easier to frequently insert and eject media from CD-ROM or floppy devices.

BZ#720691

The libvirt-guests init script attempted to make calls to the libvirtd daemon even if the daemon was inaccessible. As a consequence, the init script printed superfluous error messages that could be confusing. With this update, the script checks for a working connection, and skips calls on that connection if it is not working.

BZ#714759

This update introduces a new virsh command, "domiflist" to display detailed network interfaces information, and two new field for the "domblklist" command.

BZ#781562

Along with the "rombar" option that controls whether or not a boot ROM is made visible to the guest, QEMU also has the "romfile" option that allows specifying a binary file to present as the ROM BIOS of any emulated or pass-through PCI device. This update adds support for specifying "romfile" to both pass-through PCI devices, and emulated network devices that attach to the guest's PCI bus.

BZ#681033

Previously, libvirt did not provide means to add and display host metadata while listing guests. It was therefore impossible to store additional information about guests. A new element has been added to the libvirt XML configuration file, which allows users to store a description along with the API that allows modifications of guest metadata. The "virsh list" command has been updated to allow printing of the short description. As a result, identification of guests is now easier.

BZ#605953

This update adds a new virsh command, "iface-virsh", that allows users to "bridge" one of the host's Ethernet devices so that virtual guests can be connected directly to the physical network, rather than through a libvirt virtual network. The "iface-unbridge" command can be used to revert the interface to its previous state.

All users of libvirt are advised to upgrade to these updated packages, which fix these issues and add these enhancements.

5.167. LIBXKLAVIER

5.167.1. RHBA-2012:0923 – libxklavier bug fix update

Updated libxklavier packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The libxklavier library provides a high-level API for the X Keyboard Extension (XKB) that allows extended keyboard control. This library supports X.Org and other commercial implementations of the X Window system. The library is useful for creating XKB-related software, such as layout indicators.

Bug Fixes

BZ#657726, BZ#766645

Prior to this update, an attempt to log into the server using an NX or VNC client triggered an XInput error that was handled incorrectly by the libxklavier library due to the way how the NoMachine NX Free Edition server implements XInput support. As a consequence, the gnome-settings-daemon aborted unexpectedly. This update modifies the XInput error handling routine in the libxklavier library. Now, the library ignores this error and the gnome-settings-daemon runs as expected.

BZ#726885

Prior to this update, the keyboard layout indicator did not show if the layout was changed for the first time. As a consequence, users could, under certain circumstances, not log in. This update modifies the gnome-settings-daemon so that the indicator now shows the correct layout.

All users of libxklavier are advised to upgrade to these updated packages, which fix these bugs.

5.168. LIBXML2

5.168.1. RHSA-2012:1512 – Important: libxml2 security update

Updated libxml2 packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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(s) associated with each description below.

The libxml2 library is a development toolbox providing the implementation of various XML standards.

Security Fix

CVE-2012-5134

A heap-based buffer underflow flaw was found in the way libxml2 decoded certain entities. A remote attacker could provide a specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

All users of libxml2 are advised to upgrade to these updated packages, which contain a backported patch to correct this issue. The desktop must be restarted (log out, then log back in) for this update to take effect.

5.168.2. RHSA-2012:1288 – Moderate: libxml2 security update

Updated libxml2 packages that fix multiple security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

The libxml2 library is a development toolbox providing the implementation of various XML standards.

Security Fixes

CVE-2012-2807

Multiple integer overflow flaws, leading to heap-based buffer overflows, were found in the way libxml2 handled documents that enable entity expansion. A remote attacker could provide a large, specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2011-3102

A one byte buffer overflow was found in the way libxml2 evaluated certain parts of XML Pointer Language (XPointer) expressions. A remote attacker could provide a specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

All users of libxml2 are advised to upgrade to these updated packages, which contain backported patches to correct these issues. The desktop must be restarted (log out, then log back in) for this update to take effect.

5.169. LIBXSLT

5.169.1. RHSA-2012:1265 – Important: libxslt security update

Updated libxslt packages that fix several security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

libxslt is a library for transforming XML files into other textual formats (including HTML, plain text, and other XML representations of the underlying data) using the standard XSLT stylesheet transformation mechanism.

Security Fixes

CVE-2012-2871

A heap-based buffer overflow flaw was found in the way libxslt applied templates to nodes selected by certain namespaces. An attacker could use this flaw to create a malicious XSL file that, when used by an application linked against libxslt to perform an XSL transformation, could cause the application to crash or, possibly, execute arbitrary code with the privileges of the user running the application.

CVE-2012-2825, CVE-2012-2870, CVE-2011-3970

Several denial of service flaws were found in libxslt. An attacker could use these flaws to create a malicious XSL file that, when used by an application linked against libxslt to perform an XSL transformation, could cause the application to crash.

CVE-2011-1202

An information leak could occur if an application using libxslt processed an untrusted XPath expression, or used a malicious XSL file to perform an XSL transformation. If combined with other flaws, this leak could possibly help an attacker bypass intended memory corruption protections.

All libxslt users are advised to upgrade to these updated packages, which contain backported patches to resolve these issues. All running applications linked against libxslt must be restarted for this update to take effect.

5.170. LLDPAD

5.170.1. RHBA-2012:1175 – lldpad bug fix update

Updated lldpad packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The lldpad packages provide the Linux user space daemon and configuration tool for Intel's Link Layer Discovery Protocol (LLDP) agent with Enhanced Ethernet support.

Bug Fix

BZ#844415

Previously, an error in the DCBX (Data Center Bridging Exchange) version selection logic could cause LLDPDUs (Link Layer Discovery Protocol Data Units) to be not encoded in the TLV (Type-Length Value) format during the transition from IEEE DCBX to the legacy DCBX mode. Consequently, link flaps, a delay, or a failure in synchronizing up DCBX between the host and a peer device could occur. In the case of booting from a remote FCoE (Fibre-Channel Over Ethernet) LUN (Logical Unit Number), this bug could result in a failure to boot. This update fixes the bug and TLV is now always used in the described scenario.

All users of lldpad are advised to upgrade to these updated packages, which fix this bug.

5.170.2. RHBA-2012:1002 – lldpad bug fix update

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

The lldpad packages provide the Linux user space daemon and configuration tool for Intel's Link Layer Discovery Protocol (LLDP) agent with Enhanced Ethernet support.

Bug Fix

BZ#828684

Previously, dcbtool commands could, under certain circumstances, fail to enable the Fibre Channel over Ethernet (FCoE) application type-length-values (TLV) for a selected interface during the installation process. Consequently, various important features might have not been enabled (for

example priority flow control, or PFC) by the Data Center Bridging eXchange (DCBX) peer. To prevent such problems, application-specific parameters (such as the FCoE application TLV) in DCBX are now enabled by default.

All users of lldpad are advised to upgrade to these updated packages, which fix this bug.

5.170.3. RHBA-2012:0901 – lldpad bug fix and enhancement update

Updated lldpad packages that fix various bugs and provide an enhancement are now available for Red Hat Enterprise Linux 6.

The lldpad package provides the Link Layer Discovery Protocol (LLDP) Linux user space daemon and associated configuration tools. It supports Intel's Link Layer Discovery Protocol (LLDP) and provides Enhanced Ethernet support.

Bug Fixes

BZ#768555

The lldpad tool is initially invoked by initrd during the boot process to support Fibre Channel over Ethernet (FCoE) boot from a Storage Area Network (SAN). The runtime lldpad initscript did not kill lldpad before restarting it after system boot. Consequently, lldpad could not be started normally after system boot. In this update, lldpad init now contains the "-k" option to terminate the first instance of lldpad that was started during system boot.

BZ#803482

When the Data Center Bridging Exchange (DCBX) IEEE mode fails, it falls back to Converged Enhanced Ethernet (CEE) mode and Data Center Bridging (DCB) is enabled as part of the ifup routine. Normally, this does not occur unless either a CEE-DCBX Type-Length-Value (TLV) is received or the user explicitly enables this mode. However, in kernels released earlier than 2.6.38, DCBX IEEE mode is not supported and IEEE falls back to CEE mode immediately. Consequently, DCB was enabled in CEE mode on some kernels when IEEE mode failed, even though a peer TLV had not yet been received and the user did not manually enable it. This update fixes the logic by only enabling and advertising DCBX TLVs when a peer TLV is received. As a result, lldpad DCBX works as expected; IEEE mode is the default and CEE mode is used only if a peer CEE-DCBX TLV is received or the user enables it through the command line.

BZ#811422

A user may use dcbtool commands to clear the advertise bits on CEE-DCBX feature attributes (such as PFC, PG, APP). However, the user settings were lost during ifdown and ifup sequences and the default values were restored. This update fixes the problem so that the values are only set to defaults if the user has not explicitly enabled them.

Enhancement

BZ#812202

When a switch disassociated a connection for a virtual machine (VM) running on a host and the VM was configured to use 802.1Qbg, then libvirt was not informed and the VM connectivity was lost. Libvirt has support for restarting a VM, but it relies on the LLDP Agent Daemon to forward the Virtual Switch Interface (VSI) information. This update enables forwarding of the switch-originated VSI message to libvirt.

All users of lldpad are advised to upgrade to this updated package, which fixes these bugs and adds this enhancement.

5.171. LM_SENSORS

5.171.1. RHBA-2012:1309 – lm_sensors bug fixes

Updated lm_sensors packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The lm_sensors packages provide a set of modules for general SMBus access and hardware monitoring.

Bug Fixes

BZ#610000, BZ#623587

Prior to this update, the sensors-detect script did not detect all GenuineIntel CPUs. As a consequence, lm_sensors did not load coretemp module automatically. This update uses a more generic detection for Intel CPUs. Now, the coretemp module is loaded as expected.

BZ#768365

Prior to this update, the sensors-detect script reported an error when running without user-defined input. This behavior had no impact on the function but could confuse users. This update modifies the underlying code to allow for the sensors-detect script to run without user.

All users of lm_sensors are advised to upgrade to these updated packages, which fix these bugs.

5.172. LOGROTATE

5.172.1. RHBA-2012:1172 – logrotate bug fix update

Updated logrotate packages that fix one bug 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 Fix

BZ#827570

Attempting to send a file to a specific e-mail address failed if the "mailfirst" and "delaycompress" options were used at the same time. This was because logrotate searched for a file with the "gz" suffix, however the file had not yet been compressed. The underlying source code has been modified, and logrotate correctly finds and sends the file under these circumstances.

All users of logrotate are advised to upgrade to these updated packages, which fix this bug.

5.172.2. RHBA-2012:0786 – logrotate bug fix and enhancement update

An updated logrotate package that fixes various bugs and adds one enhancement is 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#659705

Prior to this update, the AUTHORS section of the logrotate(8) manual page contained invalid contact information for previous maintainers. This update modifies the manual page and adds this information to ensure that the contact information of the current maintainers is correct.

BZ#659173

Prior to this update, the definition of the "size" parameter in the logrotate(8) manual page was misleading. This update modifies the manual page and this definition is now easier to understand.

BZ#659720

Prior to this update, the logrotate(8) manual page did not list all options that the logrotate utility actually recognizes. This update adds the missing options and the manual list contains now all available options.

BZ#674864

Prior to this update, the logrotate(8) manual page contained several misprints. This update modifies the logrotate(8) manual page and corrects the misprints.

BZ#736053

Prior to this update, logrotate did not check the configuration file for correctly matched brackets. As a consequence, system files could be wrongly removed. This update modifies logrotate so that brackets are detected and checked for correct matching. Now, configuration files without matching brackets are skipped.

Enhancement

BZ#683622

Prior to this update, logrotate removed the Access Control Lists (ACL) flags, which are used to permit selected groups to access all logs, when rotating logs. As a consequence, selected groups could not access all logs. With this update, ACL support has been added to logrotate.

All users of logrotate are advised to upgrade to this updated package, which fixes these bugs and adds this enhancement.

5.173. LOHIT-KANNADA-FONTS

5.173.1. RHBA-2012:0870 – lohit-kannada-fonts bug fix update

An updated lohit-kannada-fonts package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The lohit-kannada-fonts package provides a free Kannada TrueType/OpenType font.

Bug Fix

BZ#603415

When using Kannada (kn_IN) locale, multiple strings could appear incomplete due to incorrect rendering of the end characters. As a consequence, users were not able to properly read the menu items on the desktop panel's menu bar. With this update, all characters are rendered correctly and users are now able to read the menu items.

All users of lohit-kannada-fonts are advised to upgrade to this updated package, which fixes this bug.

5.174. LOHIT-TELUGU-FONTS

5.174.1. RHBA-2012:1212 – lohit-telugu-fonts bug fix update

An updated lohit-telugu-fonts package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The lohit-telugu-fonts package provides a free Telugu TrueType/OpenType font.

Bug Fix

BZ#640610

Due to a bug in the lohit-telugu-fonts package, four certain syllables were rendering incorrectly. This bug has been fixed and these syllables now render correctly.

All users of lohit-telugu-fonts are advised to upgrade to this updated package, which fixes this bug.

5.175. LSOF

5.175.1. RHBA-2012:0442 – lsof bug fix update

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

The lsof (LiSt Open Files) package provides a utility to list information about files that are open and running on Linux and UNIX systems.

Bug Fixes

BZ#747375

Previously, only the first "+e" or "-e" option was processed and the rest were ignored. Consequently it was not possible to exclude more than one file system from being subjected to kernel function calls. This update fixes this issue and lsof now functions as expected with multiple +e or -e options.

BZ#795799

Prior to this update, the lsof utility ignored the "-w" option if both the "-b" and the "-w" options were specified. As a consequence, lsof failed to suppress warning messages. Now, the -w option successfully suppresses warning messages.

All users of lsof are advised to upgrade to this updated package, which fixes these bugs.

5.176. LSVPD

5.176.1. RHBA-2012:0795 – lsvpd bug fix update

Updated lsvpd packages that fix four bugs are now available for Red Hat Enterprise Linux 6.

The lsvpd package provides the lsvpd command to list device Vital Product Data (VPD), the lscfg command to display configuration, diagnostic, and vital product data (VPD) information, and the lsmcode command to display the microcode and firmware levels.

Bug Fixes

BZ#684646

Prior to this update, the vpdupdate tool tried to collect additional information about the SCSI devices. As a consequence, vpdupdate sent irrelevant messages to the syslog file. This update modifies the sysfs layout in lsvpd. Now, no more unwanted messages are sent to the syslog file.

BZ#688574

Prior to this update, the lsvpd man page did not correctly describe the "-p" option. With this update, the man page correctly states that the "-p" option prints the designed output.

BZ#714086

Prior to this update, the lscfg(8), lsmcode(8), lsmsr(8), lsvio(8), and vpdupdate(8) man pages missed to document several lsvpd options. This update adds all missing options to the man pages.

BZ#741899

Prior to this update, the lsmcode tool did not contain the build requirement "librtas-devel". As a consequence, the firmware version was not displayed. This update modifies lsmcode file so that the missing build requirement is added to the spec file.

All users of lsvpd are advised to upgrade to this updated package, which adds this enhancement.

5.177. LTRACE

5.177.1. RHBA-2012:0926 – ltrace bug fix and enhancement update

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

The ltrace utility is a debugging program that runs a specified command until the command exits. While the command is executing, ltrace intercepts and records both the dynamic library calls called by the executed process and the signals received by the executed process. The ltrace utility can also intercept and print system calls executed by the process.

Bug Fixes

BZ#742340

Prior to this update, a traced process that had more than one thread could be aborted if the threads ran into breakpoints which the ltrace utility did not handle. With this update, ltrace attaches to the newly created threads, and carefully handles the breakpoints so that tracing events are not missed. This update also improves the detach logic so that a running process to which ltrace has been attached is left in a consistent state before detaching.

BZ#811184

Prior to this update, the ltrace utility could, under certain circumstances, fail to trace returns from functions which where called with a tail call optimization. This update adds support for tracing returns from functions called with a tail call optimization.

Enhancement

BZ#738254

Prior to this update, ltrace could not trace library functions loaded via libdl. This update changes the behavior of the "-x" option for placing static breakpoints so that dynamic libraries are also considered and breakpoints set in them. This works for dynamic libraries that are linked to the binary as well as those that are opened with the function "dlopen" in runtime.

All users of ltrace are advised to upgrade to this updated package, which fixes these bugs and adds this update.

5.178. LUCI

5.178.1. RHBA-2012:0766 – luci bug fix and enhancement update

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

The luci packages contain a web-based high-availability cluster configuration application.

Bug Fixes

BZ#796731

A cluster configuration can define global resources (declared outside of cluster service groups) and in-line resources (declared inside a service group). The names of resources must be unique regardless of whether the resource is a global or an in-line resource. Previously, **luci** allowed a global resource with a name, which was already used by a resource that had been declared in-line, and a name of a resource within a service group with a name that was already used by a global resource. As a result, luci could terminate unexpectedly with error 500 or the cluster configuration could be modified improperly. With this update, luci fails gracefully under these circumstances and reports the problems, and the cluster configuration remains unmodified.

BZ#749668

If a cluster configuration was not valid, **luci** terminated unexpectedly without any report about the cluster configuration problem. It was therefore impossible to use luci for administration of clusters with invalid configurations. With this update, luci detects invalid configurations and returns warnings with information about possible mistakes along with proposed fixes.

BZ#690621

The user could not debug problems on-the-fly if debugging was not enabled prior to starting **luci**. This update adds new controls to allow the user to change the log level of messages generated by luci according to the message type while luci is running.

BZ#801491

When the user created a cluster resource with a name that contained a period symbol (**.**), **luci** failed

to redirect the browser to the resource that was just created. As a result, error 500 was displayed, even though the resource was created correctly. This update corrects the code that handles redirection of the browser after creating such a resource and luci redirects the browser to a screen that displays the resource as expected.

BZ#744048

Previously, **luci** did not require any confirmation on removal of cluster services. Consequently, the user could remove the services by accident or without properly considering the consequences. With this update, luci displays a confirmation dialog when the user requests removal of cluster services, which informs the user about the consequences, and forces them to confirm their action.

BZ#733753

Since **Red Hat Enterprise Linux 6.3**, authenticated sessions automatically expired after 15 minutes of inactivity. With this update, the user can now change the time-out period in the who.auth_tkt_timeout parameter in the **/etc/sysconfig/luci** file.

BZ#768406

Previously, the default value of the monitor_link attribute of the IP resource agent was displayed incorrectly: when not specified explicitly, its value was displayed as enabled while it was actually disabled, and vice versa. When the user made changes to the monitor_link value using **luci**, an incorrect value was stored. With this update, the monitor_link value is display properly, and the user can now view and modify the value as expected.

BZ#755092

The **force_unmount** option was not shown for file-system resources and the user could not change the configuration to enable or disable this option. A checkbox that displays its current state was added and the user can now view and change the force_unmount attribute of file-system resources.

BZ#800239

A new attribute, **tunneled**, was added to the VM (Virtual Machine) resource agent script. This update adds a checkbox displaying the current value of the **tunneled** attribute to the VM configuration screen so that the user can enable or disable the attribute.

BZ#772314

Previously, an ACL (Access Control List) system was added to allow delegation of permissions to other **luci** users. However, permissions could not be set for users until they had logged in at least once. With this update, ACLs can be added and changed before the user logs in to **luci** for the first time.

BZ#820402

Due to a regression, the Intel Modular and IF MIB fencing agents were removed from the list of devices for which users could configure new instances. Consequently, users could not create a new instance of these fencing devices. The Intel Modular and IF MIB fencing device entries have been added back to the list of fence devices and users are again able to create new instances of Intel Modular and IF MIB fencing devices.

Enhancements

BZ#704978

In the **Create and edit service groups** form, the relationships between groups could not always be

easily discerned. Solid borders were added along the side of resources within the forms to make the relationships between resources clearer. Also, when adding a resource to a service group, the screen is scrolled to the resource that was added.

BZ#740835

While creating and editing fail-over domains, the user could select and unselect checkboxes and enter values into text fields whose values were ignored. Such checkboxes and text fields are now disabled and become enabled only when their values are used.

BZ#758821

To provide an interface for working with the RRP (Redundant Ring Protocol) configuration in **luci**, Technology Preview support was added for RRP in the corosync cluster engine that the Red Hat HA stack is built upon. The **Redundant Ring** configuration tab is now available in the **Configure** tab of clusters to allow RRP configuration from luci.

BZ#786584

A new resource agent was added to provide high-availability of condor-related system daemons. This update adds support for viewing, creating, and editing the configuration of Condor resources to allow the user to configure the Condor resource agent.

BZ#707471

The **reboot** icon was similar to the **refresh** icon and the user could have mistakenly rebooted a cluster node instead of refreshing the status information. With this update, the **reboot** icon has been changed. Also, a dialog box is now displayed before reboot so the user must confirm their reboot request.

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

5.179. LVM2

5.179.1. RHEA-2012:1574 – lvm2 enhancement update

Updated lvm2 packages that add an enhancement are now available for Red Hat Enterprise Linux 6.

The lvm2 packages provide support for Logical Volume Management (LVM).

Enhancement

BZ#883034

In cases of transient inaccessibility of a PV (Physical Volume), such as with iSCSI or other unreliable transport, LVM required manual action to restore the PV for use even if there was no room for conflict. With this update, the manual action is no longer required if the transiently inaccessible PV had no active metadata areas (MDA). The automatic restore action of a physical volume (PV) from the MISSING state after it becomes reachable again and if it has no active MDA has been added to the lvm2 packages.

Users of lvm2 are advised to upgrade to these updated packages, which adds this enhancement.

5.179.2. RHBA-2012:1399 – lvm2 bug fix update

Updated lvm2 packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The lvm2 packages provide support for Logical Volume Management (LVM).

Bug Fixes

BZ#843808

When using a physical volume (PV) that contained ignored metadata areas, an LVM command, such as pvs, could incorrectly display the PV as being orphan despite it belonged to a volume group (VG). This incorrect behavior was also dependent on the order of processing each PV in the VG. With this update, the processing of PVs in a VG has been fixed to properly account for PVs with ignored metadata areas so that the order of processing is no longer important, and LVM commands now always give the same correct result, regardless of PVs with ignored metadata areas.

BZ#852438

Previously, if the "issue_discards=1" configuration option was used with an LVM command, moving physical volumes using the pvmove command resulted in data loss. This update fixes the bug in pvmove and the data loss no longer occurs in the described scenario.

BZ#852440

When the "--alloc anywhere" command-line option was specified for the lvcreate command, an attempt to create a logical volume failed if "raid4", "raid5", or "raid6" was specified for the "--type" command-line option as well. A patch has been provided to address this bug and lvcreate now succeeds in the described scenario.

BZ#852441

An error in the way RAID 4/5/6 space was calculated, was preventing users from being able to increase the size of these logical volumes. This update provides a patch to fix this bug but it comes with two limitations. Firstly, a RAID 4/5/6 logical volume cannot be reduced in size yet. Secondly, users cannot extend a RAID 4/5/6 logical volume with a different stripe count than the original.

BZ#867009

If the "issue_discards=1" configuration option was set in the /etc/lvm/lvm.conf file, it was possible to issue a discard request to a PV that was missing in a VG. Consequently, the dmeventd, lvremove, or vgreduce utilities could terminate unexpectedly with a segmentation fault. This bug has been fixed and discard requests are no longer issued on missing devices. As the discard operation is irreversible, in addition to this fix, a confirmation prompt has been added to the lvremove utility to ask the user before discarding a LV, thus increasing robustness of the discard logic.

Users of lvm2 are advised to upgrade to these updated packages, which fix these bugs.

5.179.3. RHBA-2012:0962 – lvm2 bug fix and 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 contain support for *Logical Volume Management* (LVM).

Bug Fixes

BZ#683270

When mirrors are up-converted it was impossible to add another image (copy) to a mirrored logical volume whose activation was regulated by tags present in the ***volume_list*** parameter of **lvm.conf**. The code has been improved to temporarily copy the mirror's tags to the in-coming image so that it can be properly activated. As a result, high-availability (HA) **LVM** service relocation now works as expected.

BZ#700128

Previously, the **lvremove** command could fail with the error message "Can't remove open logical volume" despite the volume itself not being in use anymore. In most cases, such a situation was caused by the udev daemon that still processed events that preceded the removal and it kept the device open while lvremove tried to remove it at the same time. This update adds a retry loop to help to avoid this problem. The removal is tried several times before the command fails completely.

BZ#733522

Previously, if a snapshot with a virtual origin was created in a clustered *Volume Group* (VG), it incorrectly tried to activate on other nodes as well and the command failed with "Error locking on node" error messages. This has been fixed and a snapshot with virtual origin, using **--virtualsize**, is now properly activated exclusively only (on local node).

BZ#738484

Previously, if **clvmd** received an invalid request through its socket (for example an incomplete header was sent), the clvmd process could terminate unexpectedly or stay in an infinite loop. Additional checks have been added so that such invalid packets now cause a proper error reply to the client and clvmd no longer crashes in the scenario described.

BZ#739190

The device-mapper daemon (**dmeventd**) is used, for example, for monitoring **LVM** based mirrors and snapshots. When attempting to create a snapshot using **lvm2**, the **lvcreate -s** command resulted in a **dlopen** error if dmeventd was upgraded after the last system restart. With this update, dmeventd is now restarted during a package update to fetch new versions of installed libraries to avoid any code divergence that could end up with a symbol lookup failure.

BZ#740290

Restarting **clvmd** with option **-S** should preserve exclusive locks on a restarted cluster node. However the option **-E**, which should pass such exclusive locks, had errors in its implementation. Consequently, exclusive locks were not preserved in a cluster after restart. This update implements proper support for option **-E**. As a result, after restarting clvmd the locks will preserve a cluster's exclusive state.

BZ#742607

If a device-mapper device was left open by a process, it could not be removed with the **dmsetup [--force] remove device_name** command. The **--force** option failed, reporting that the device was busy. Consequently, the underlying block device could not be detached from the system. With this update, dmsetup has a new command **wipe_table** to wipe the table of the device. Any subsequent I/O sent to the device returns errors and any devices used by the table, that is to say devices to which the I/O is forwarded, are closed. As a result, if a long-running process keeps a device open after it has finished using it, the underlying devices can be released before that process exits.

BZ#760946

Using a prefix or command names on **LVM** command output (the **log/prefix** and **log/command_names** directive in **lvm.conf**) caused the **lvm2-monitor** init script to fail to start monitoring for relevant VGs. The init script acquires the list of VGs first by calling the **vgs** command

and then it uses its output for further processing. However, if the prefix or command name directive is used on output, the VG name was not correctly formatted. To solve this, the lvm2-monitor init script now overrides the log/prefix and log/command_names setting so the command's output is always suitable for use in the init script.

BZ#761267

Prior to this update, the **lvconvert --merge** command did not check if the snapshot in question was invalid before proceeding. Consequently, the operation failed part-way through, leaving an invalid snapshot. This update disallows an invalid snapshot to be merged. In addition, it allows the removal of an invalid snapshot that was to be merged on next activation, or that was invalidated while merging (the user will be prompted for confirmation).

BZ#796602

If a Volume Group name is supplied together with the Logical Volume name, **lvconvert --splitmirrors** fails to strip it off. This leads to an attempt to use a Logical Volume name that is invalid. This release detects and validates any supplied Volume Group name correctly.

BZ#799071

Previously, if **pvmove** was used on a clustered VG, temporarily activated pvmove devices were improperly activated cluster-wide (that is to say, on all nodes). Consequently, in some situations, such as when using tags or HA-LVM configuration, pvmove failed. This update fixes the problem, pvmove now activates all such devices exclusively if the Logical Volumes to be moved are already exclusively activated.

BZ#807441

Previously, when the **vgreduce** command was executed with a non-existent VG, it unnecessarily tried to unlock the VG at command exit. However the VG was not locked at that time as it was unlocked as part of the error handling process. Consequently, the vgreduce command failed with the internal error "Attempt to unlock unlocked VG " when it was executed with a non-existent VG. This update improves the code to provide a proper check so that only locked VGs are unlocked at vgreduce command exit.

BZ#816711

Previously, requests for information regarding the health of the log device were processed locally. Consequently, in the event of a device failure that affected the log of a cluster mirror, it was possible for a failure to be ignored and this could cause I/O to the mirror LV to become unresponsive. With this update, the information is requested from the cluster so that log device failures are detected and processed as expected.

Enhancements

BZ#464877

Most **LVM** commands require an accurate view of the LVM metadata stored on the disk devices on the system. With the current LVM design, if this information is not available, LVM must scan all the physical disk devices in the system. This requires a significant amount of I/O operations in systems that have a large number of disks. The purpose of the LV Metadata daemon (**lvmetad**) is to eliminate the need for this scanning by dynamically aggregating metadata information each time the status of a device changes. These events are signaled to **lvmetad** by udev rules. If **lvmetad** is not running, LVM performs a scan as it normally would. This feature is provided as a Technology Preview and is

disabled by default in Red Hat Enterprise Linux 6.3. To enable it, refer to the ***use_lvmetad*** parameter in the **/etc/lvm/lvm.conf** file, and enable the **lvmetad** daemon by configuring the lvm2-lvmetad init script.

BZ#593119

The expanded RAID support in **LVM** is now fully supported in Red Hat Enterprise Linux 6.3, with the exception of RAID logical volumes in HA-LVM. LVM now has the capability to create RAID 4/5/6 Logical Volumes and supports a new implementation of mirroring. The MD (software RAID) modules provide the back-end support for these new features.

BZ#637693

When a new LV is defined, anyone who has access to the LV can read any data already present on the LUNs in the extents allocated to the new LV. Users can create thin volumes by using the "lvcreate -T" command to meet the requirement that zeros are returned when attempting to read a block not previously written to. The default behavior of thin volumes is the provisioning of zero data blocks. The size of provisioned blocks is in the range of 64KB to 1GB. The bigger the blocks the longer it takes for the initial provisioning. After first write, the performance should be close to a native linear volume. However, for a clustering environment there is a difference as thin volumes may be only exclusively activated.

BZ#658639

This update greatly reduces the time spent on creating identical data structures, which allows even a very large number of devices (in the thousands) to be activated and deactivated in a matter of seconds. In addition, the number of system calls from device scanning has been reduced, which also gives a 10%-30% speed improvement.

BZ#672314

Some **LVM** segment types, such as "mirror", have single machine and cluster-aware variants. Others, such as snapshot and the RAID types, have only single machine variants. When switching the cluster attribute of a Volume Group (VG), the aforementioned segment types must be inactive. This allows for re-loading of the appropriate single machine or cluster variant, or for the necessity of the activation to be exclusive in nature. This update disallows changing cluster attributes of a VG while RAID LVs are active.

BZ#731785

The **dmsetup** command now supports displaying block device names for any devices listed in the "deps", "ls" and "info" command output. For the dmsetup "deps" and "ls" command, it is possible to switch among "devno" (major and minor number, the default and the original behavior), "devname" (mapping name for a device-mapper device, block device name otherwise) and "blkdevname" (always display a block device name). For the dmsetup "info" command, it is possible to use the new "blkdevname" and "blkdevs_used" fields.

BZ#736486

Device-mapper allows any character except "/" to be used in a device-mapper name. However, this is in conflict with udev as its character whitelist is restricted to 0-9, A-Z, a-z and #+-.:=@_. Using any black-listed character in the device-mapper name ends up with incorrect **/dev** entries being created by udev. To solve this problem, the **libdevmapper** library together with the **dmsetup** command now supports encoding of udev-blacklisted characters by using the "\xNN" format where NN is the hex value of the character. This format is supported by udev. There are three "mangling" modes in which libdevmapper can operate: "none" (no mangling), "hex" (always mangle any blacklisted character) and "auto" (use detection and mangle only if not mangled yet). The default mode used is "auto" and any libdevmapper user is affected unless this setting is changed by the respective libdevmapper call.

To support this feature, the dmsetup command has a new ***--manglename <mangling_mode>*** option to define the name mangling mode used while processing device-mapper names. The ***dmsetup info -c -o*** command has new fields to display: "mangled_name" and "unmangled_name". There is also a new **dmsetup mangle** command that renames any existing device-mapper names to its correct form automatically. It is strongly advised to issue this command after an update to correct any existing device-mapper names.

BZ#743640

It is now possible to extend a mirrored logical volume without inducing a synchronization of the new portion. The "--nosync" option to **lvextend** will cause the initial synchronization to be skipped. This can save time and is acceptable if the user does not intend to read what they have not written.

BZ#746792

LVM mirroring has a variety of options for the bitmap write-intent log: "core", "disk", "mirrored". The cluster log daemon (cmirrord) is not multi-threaded and can handle only one request at a time. When a log is stacked on top of a mirror (which itself contains a 'core' log), it creates a situation that cannot be solved without threading. When the top level mirror issues a "resume", the log daemon attempts to read from the log device to retrieve the log state. However, the log is a mirror which, before issuing the read, attempts to determine the "sync" status of the region of the mirror which is to be read. This sync status request cannot be completed by the daemon because it is blocked on a read I/O to the very mirror requesting the sync status. With this update, the "mirrored" option is not available in the cluster context to prevent this problem from occurring.

BZ#769293

A new **LVM** configuration file parameter, ***activation/read_only_volume_list***, makes it possible to activate particular volumes always in read-only mode, regardless of the actual permissions on the volumes concerned. This parameter overrides the **--permission rw** option stored in the metadata.

BZ#771419

In previous versions, when monitoring of multiple snapshots was enabled, **dmeventd** would log redundant informative messages in the form "Another thread is handling an event. Waiting... ". This needlessly flooded system log files. This behavior has been fixed in this update.

BZ#773482

A new implementation of **LVM** copy-on-write (cow) snapshots is available in Red Hat Enterprise Linux 6.3 as a Technology Preview. The main advantage of this implementation, compared to the previous implementation of snapshots, is that it allows many virtual devices to be stored on the same data volume. This implementation also provides support for arbitrary depth of recursive snapshots. This feature is for use on a single-system. It is not available for multi-system access in cluster environments. For more information, refer to the documentation of the **-s** or **--snapshot** option in the **lvcreate** man page.

BZ#773507

Logical Volumes (LVs) can now be thinly provisioned to manage a storage pool of free space to be allocated to an arbitrary number of devices when needed by applications. This allows creation of devices that can be bound to a thinly provisioned pool for late allocation when an application actually writes to the LV. The thinly-provisioned pool can be expanded dynamically if and when needed for cost-effective allocation of storage space. In Red Hat Enterprise Linux 6.3, this feature is introduced as a Technology Preview. For more information, refer to the lvcreate man page. Note that the device-mapper-persistent-data package is required.

BZ#796408

LVM now recognizes EMC PowerPath devices (emcpower) and uses them in preference to the devices out of which they are constructed.

BZ#817130

LVM now has two implementations for creating mirrored logical volumes: the "mirror" segment type and the "raid1" segment type. The "raid1" segment type contains design improvements over the "mirror" segment type that are useful to its operation with snapshots. As a result, users who employ snapshots of mirrored volumes are encouraged to use the "raid1" segment type rather than the "mirror" segment type. Users who continue to use the "mirror" segment type as the origin LV for snapshots should plan for the possibility of the following disruptions.

When a snapshot is created or resized, it forces I/O through the underlying origin. The operation will not complete until this occurs. If a device failure occurs to a mirrored logical volume (of "mirror" segment type) that is the origin of the snapshot being created or resized, it will delay I/O until it is reconfigured. The mirror cannot be reconfigured until the snapshot operation completes, but the snapshot operation cannot complete unless the mirror releases the I/O. Again, the problem can manifest itself when the mirror suffers a failure simultaneously with a snapshot creation or resize.

There is no current solution to this problem beyond converting the mirror from the "mirror" segment type to the "raid1" segment type. In order to convert an existing mirror from the "mirror" segment type to the "raid1" segment type, perform the following action:

```
~]$ lvconvert --type raid1 <VG>/<mirrored LV>
```

This operation can only be undone using the **vgcfgrestore** command.

With the current version of **LVM2**, if the "mirror" segment type is used to create a new mirror LV, a warning message is issued to the user about possible problems and it suggests using the "raid1" segment type instead.

Users of lvm2 should upgrade to these updated packages, which fix these bugs and add these enhancements.

5.179.4. RHBA-2013:1472 – lvm2 bug fix update

Updated lvm2 packages that fix one bug 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 Fix

BZ#965810

Previously, on certain HP servers using Red Hat Enterprise Linux 6 with the xfs file system, a regression in the code caused the lvm2 utility to ignore the "optimal_io_size" parameter and use a 1MB offset start. Consequently, there was an increase in the disk write operations which caused data misalignment and considerably lowered the performance of the servers. With this update, lvm2 no longer ignores "optimal_io_size" and data misalignment no longer occurs in this scenario.

Users of lvm2 are advised to upgrade to these updated packages, which fix this bug.

5.180. M2CRYPTO

5.180.1. RHBA-2012:0975 – m2crypto bug fix update

An updated m2crypto package that fixes various bugs is now available for Red Hat Enterprise Linux 6.

The m2crypto package allows Python programs to call OpenSSL functions.

Bug Fixes

BZ#742914

The M2Crypto.httpslib.HTTPSConnection class always created an IPv4 socket. This made it impossible to connect to IPv6 servers using this class. With this update, the implementation now correctly creates an IPv4 or IPv6 socket, as necessary, thus adding support for IPv6 servers.

BZ#803520

Prior to this update, the AES_crypt() function did not free a temporary buffer. This caused a memory leak when the function was called repeatedly. This problem has been fixed and the AES_crypt() function now frees memory correctly.

BZ#803554

The implementation of HTTPS connections via a proxy did not reflect the changes in Python 2.6. Consequently, every attempt to connect to an HTTPS server using a proxy failed, generating an exception. The M2Crypto implementation has been updated and now works correctly.

All users of m2crypto are advised to upgrade to this updated package, which fixes these bugs.

5.181. MAILMAN

5.181.1. RHBA-2012:1474 – mailman bug fix update

Updated mailman packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

Mailman is a program used to help manage e-mail discussion lists.

Bug Fixes

BZ#772998

The reset_pw.py script contained a typo, which could cause the mailman utility to fail with a traceback. The typo has been corrected, and mailman now works as expected.

BZ#799323

The "urlhost" argument was not handled in the newlist script. When running the "newlist" command with the "--urlhost" argument specified, the contents of the index archive page was not created using proper URLs; the hostname was used instead. With this update, "urlhost" is now handled in the newlist script. If the "--urlhost" argument is specified on the command line, the host URL is used when creating the index archive page instead of the hostname.

BZ#832920

Previously, long lines in e-mails were not wrapped in the web archive, sometimes requiring excessive horizontal scrolling. The "white-space: pre-wrap;" CSS style has been added to all templates, so that long lines are now wrapped in browsers that support that style.

BZ#834023

The "From" string in the e-mail body was not escaped properly. A message containing the "From" string at the beginning of a line was split and displayed in the web archive as two or more messages. The "From" string is now correctly escaped, and messages are no longer split in the described scenario.

All users of mailman are advised to upgrade to these updated packages, which fix these bugs.

5.182. MAKE

5.182.1. RHBA-2012:0443 – make bug fix update

An updated make package that fixes one bug is now available for Red Hat Enterprise Linux 6.

GNU make is a tool for controlling the generation of executables and other non-source files of a program from the program's source files. Users can build and install packages by using make without any significant knowledge about the details of the build process.

Bug Fix

BZ#699911

Prior to this update, memory corruption could occur in an "eval" expression if one of its sub-expressions was assigned to the same variable. An upstream patch has been applied to address this issue, and memory corruption no longer appears in the described scenario.

All users of make are advised to upgrade to this updated package, which fixes this bug.

5.183. MAN-PAGES-FR

5.183.1. RHBA-2012:0463 – man-pages-fr bug fix update

An updated man-pages-fr package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The man-pages-fr package contains a collection of manual pages from the man-pages Project, translated into French. It also includes supplemental pages provided by Dr. Patrick Atlas and Dr. Gerard Delafond.

Bug Fix

BZ#613622

Prior to this update, the mansupfr.tar.bz2 tarball that contains supplemental French manual pages, was not correctly extracted. As a consequence, some manual pages were not installed. With this update, supplemental manual pages are added back when no known file conflicts appear, and the supplemental French manual pages are again available.

All users, requiring localized manual pages in French, are advised to upgrade to this updated package, which fixes this bug.

5.184. MAN-PAGES-OVERRIDES

5.184.1. RHBA-2012:0961 – man-pages-overrides bug fix update

Updated man-pages-overrides packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The man-pages-overrides package provides a collection of manual (man) pages to complement other packages or update those contained therein.

Bug Fixes

BZ#528879

Prior to this update, manual pages for the mksquashfs and unsquashfs utilities were missing. This update adds the mksquashf(1) and unsquashfs(1) manual pages.

BZ#529335

Prior to this update, manual pages for the FUSE utilities fusermount, fuse, and ulockmgr_server were missing. This update adds the fusermount(1), mount.fuse(8), and ulockmgr_server(1) manual pages to the man-pages-overrides packages.

BZ#605521

Prior to this update, a manual page for the urlgrabber utility was missing. This update adds the urlgrabber(1) manual page to the man-pages-overrides packages.

BZ#653908

Prior to this update, the replcon(1) manual page was missing the description of the "-R" and "--regex" options. This update adds this description to the man-pages-overrides packages.

BZ#695363, BZ#801783

Prior to this update, the iptables(8), ip6tables(8) and ebtables(8) manual pages were missing a description of the AUDIT target module. This update adds a description of this module to the these manual pages.

BZ#745467

Prior to this update, a manual page for the pkcs_slot utility was missing. This update adds the pkcs_slot(1) manual page to the man-pages-overrides packages.

BZ#747970

Prior to this update, the lsblk(8) manual page was missing the description for the "-D" option. This update adds this description.

BZ#768949

Prior to this update, the manual page for trap in Bash did not mention that signals ignored upon entry cannot be listed later. This update modifies the text to mention that "Signals ignored upon entry to the shell cannot be trapped, reset or listed".

BZ#766341

Prior to this update, the cgcreate(1) manual page contained the invalid "-s" option in the synopsis. This update removes this option.

BZ#769566

Prior to this update, the wbinfo(1) manual page contained an incorrect description of the "--group-info" option. This update modifies this description.

BZ#800256

Prior to this update, the shmat(2) manual page was missing the description for the EIDRM error. With this update, this description is added to the shmat(2) manual page.

BZ#800385

The manual pages expect, logrotate2, logrotate3, logrotate4, logrotate5, logrotate, pcre2, pcre, lsvpd, nfs-utils, and vsftpd have been fixed in their original packages. This update removes these manual pages from the man-pages-overrides packages.

BZ#801742

Prior to this update, the request-key.conf(5) manual page contained a misprint in one sentence. With this update, this misprint is removed.

BZ#801784

Prior to this update, the yum(8) manual page contained a misprint. With this update, this misprint is corrected.

BZ#810910

Prior to this update, the mount(8) man page did not include the default option "relatime". This update includes this option in the list of options.

All users of man-pages-overrides are advised to upgrade to this updated package, which fixes these bugs.

5.185. MAN

5.185.1. RHBA-2012:0449 – man bug fix update

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

The man package provides the man, apropos, and whatis tools for finding information and documentation about the Linux system.

Bug Fixes

BZ#659646

Previously, the Japanese version of the man(1) manual page contained a duplicate line in the specification of the "-p pager" option. This update removes the duplicate.

BZ#749290

Prior to this update, the makewhatis script, which creates the whatis database of manual pages, ignored symbolic links between pages. With this update, the makewhatis script includes symbolic links in the whatis database.

All users of man are advised to upgrade to this updated package, which fixes these bugs.

5.186. MATAHARI

5.186.1. RHBA-2012:0844 – matahari bug fix and enhancement update

Updated matahari packages that fix multiple bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The matahari packages provide a set of APIs for operating system management that are exposed to remote access over the Qpid Management Framework (QMF).



NOTE

The Matahari agent framework (matahari-*) packages are deprecated starting with the Red Hat Enterprise Linux 6.3 release. 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. It is strongly recommended that users discontinue the use of the matahari packages and other packages which depend on the Matahari infrastructure (specifically, libvirt-qmf and fence-virtd-libvirt-qpid). It is recommended that users uninstall Matahari from their systems to remove any possibility of security issues being exposed.

Users who choose to continue to use the Matahari agents should note the following:

- The matahari packages are not installed by default starting with Red Hat Enterprise Linux 6.3 and are not enabled by default to start on boot when they are installed. Manual action is needed to both install and enable the matahari services.
- The default configuration for qpid (the transport agent used by Matahari) does not enable access control lists (ACLs) or SSL. Without ACLs/SSL, the Matahari infrastructure is not secure. Configuring Matahari without ACLs/SSL is not recommended and may reduce your system's security.
- The matahari-services agent is specifically designed to allow remote manipulation of services (start, stop). Granting a user access to Matahari services is equivalent to providing a remote user with root access. Using Matahari agents should be treated as equivalent to providing remote root SSH access to a host.
- By default in Red Hat Enterprise Linux, the Matahari broker (qpidd running on port 49000) does not require authentication. However, the Matahari broker is not remotely accessible unless the firewall is disabled, or a rule is added to make it accessible. Given the capabilities exposed by Matahari agents, if Matahari is enabled, system administrators should be extremely cautious with the options that affect remote access to Matahari.

Note that Matahari will not be shipped in future releases of Red Hat Enterprise Linux (including Red Hat Enterprise Linux 7), and may be considered for formal removal in a future release of Red Hat Enterprise Linux 6.

Bub Fix

BZ#752325

Prior to this update, matahari agents were being unnecessarily restarted during upgrades. As a consequence, unexpected output could appear during the upgrade process. This update modifies the underlying code so that agents are not restarted more than once and no more unexpected

reporting occurs.

Enhancements

BZ#723078

Prior to this update, no shell tool for using matahari agents was available. This update adds a Python API and command-line shell to matahari.

BZ#759243

Prior to this update, the matahari interface could not identify Python scripts written by users. This update adds the RPC agent to provide an API to execute user-written Python scripts installed on the target machine.

All users of matahari are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.187. MCELOG

5.187.1. RHBA-2012:0779 – mcelog bug fix and enhancement update

Updated mcelog packages that fix three bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The mcelog packages provide the mcelog daemon to collect and decode Machine Check Exception data on AMD64 and Intel 64 platforms.

Bug Fixes

BZ#728265

Prior to this update, the mcelog README file contained references to nonexistent directories. This update removes these references and updates the file.

BZ#769363

Prior to this update, the mcelog daemon wrongly displayed an error that a certain microarchitecture was not supported even if the CPU was supported if mcelog was run on Intel CPUs only with architectural decoding enabled. This update removes this message.

BZ#784091

Prior to this update, a cron job tried to install regardless whether a system was supported or not. As a result, the mcelog daemon displayed the message "No such device" if mcelog was installed on unsupported systems. This update prevents the cron job from installing on unsupported processors.

Enhancements

BZ#746785

Prior to this update, the mcelog daemon displayed the error "mcelog read: No such device" when running the unsupported AMD Family 16 microarchitecture or higher. This update adds a check to mcelog to determine what AMD processor family is used. If needed, the new message "CPU is unsupported" is displayed.

BZ#795508

Prior to this update, The cron file for mcelog did not use the "--supported" option. As a consequence, the "--supported" option did not correctly check whether the mcelog daemon worked. This update adds the "--supported" option to the crontab file and removes two redundant strings.

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

5.188. MDADM

5.188.1. RHBA-2012:0787 – mdadm bug fix and enhancement update

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

The mdadm package contains a utility for creating, managing, and monitoring Linux MD (multiple disk) devices.



NOTE

The mdadm package has been upgraded to upstream version 3.2.3, which provides a number of bug fixes and enhancements over the previous version. (BZ#745802)

Bug Fixes

BZ#771332

Previously, removing a device from its container occasionally failed. With this update, mdadm attempts to remove such a device again.

BZ#808776

The **mdadm --add** command could fail when adding to an array a device similar to a recent member of the array. With this update, the code restricting device addition has been corrected to only apply to members of recent arrays which have failed and require manual assembly.

BZ#811282

It was not possible to add a write-intent bitmap to a multiple-device (MD) array built with the metadata version property **1.0**. This happened because **mdadm** occasionally failed to calculate the bitmap location correctly. With this update, a write-intent bitmap can be added to such an MD array as expected.

BZ#730052

If the user rebooted the system during a reshape process, MD RAID devices could remain inactive due to several problems. This update introduces several patches adjusting the logic of the reshape process and fixing several errors.

BZ#808424

The **mdadm --monitor** command terminated unexpectedly after completing a resynchronization process due to buffer overflowing. This happened when there were more than 40 mismatches reported during the resync process as the respective buffer could hold only 40 mismatch reports. With this update, the buffer has been enlarged and can now hold up to 80 mismatch reports.

BZ#790394

The mdadm tool could fail to add a device to a degraded array with bitmaps and exited silently. This happened because the called functions attempted to write to not-aligned buffers. With this update, an upstream patch to fix this bug has been applied: the patch modifies the underlying functions to use their own aligned buffers and the problem no longer occurs.

BZ#788022

If the user installed Red Hat Enterprise Linux 6.0 on a machine with a version 0.90 MD RAID array already installed, the array did not automatically start at the next reboot. This happened because the policy statement in the **/etc/mdadm.conf** file excluded automatic startup of version 0.90 RAID arrays (however, if such an array was needed on boot, dracut did start this array). The user could add **+0.90** to the **AUTO** line in the **/etc/mdadm.conf** file to have such RAID arrays come up on boot.

In Red Hat Enterprise Linux 6.1 and 6.2, mdadm always assembled version 0.90 RAID arrays automatically due to a bug. This update fixes the bug. The user needs to implement the same fix as described above for Red Hat Enterprise Linux 6.0 to have version 0.90 RAID arrays come up on boot (add **+0.90** to the **AUTO** line in **/etc/mdadm.conf**).

BZ#808438

The **mdadm** utility did not check how many volumes per controller were allowed for arrays on Intel controllers with OROM (Option ROM) or EFI (Extensible Firmware Interface) created by IMSM (Intel Matrix Storage Manager). Consequently, only the respective limited number of arrays were available even if further arrays were previously added. With this update, **mdadm** checks OROM limitations and adding a new volume is blocked if the number of volumes on the device attached to the given controller has exceeded the limit.

BZ#771554

The **mdadm** tool did not apply the **--oneshot/-1** option when running the **mdadm --monitor --scan --oneshot** command or its short equivalent. Consequenlty, mdadm was monitoring the respective device continuously. With this update, the underlying code has been modified and the **--oneshot** option is applied as expected.

BZ#808492

IMSM RAID only supports two volumes per container. Previously, mdadm allowed an administrator to create the second volume smaller than the remaining free space leaving unallocated space in the container. With this update, the second volume must take up the remaining space of the container by applying the same restrictions as when managing the IMSM RAID throught the BIOS interface.

BZ#808507

Disk and volume sizes were stored as metadata in 32-bit blocks. If the disk size exceeded 2 TB, the allocated space was no longer sufficient to store the size value and volume sizes returned incorrect results for such disks. With this update, the size of the block is calculated depending on the disk or volume size, and, in the scenario described, the correct disk and volume sizes are used and displayed.

BZ#808519

The **mdadm** tool failed to create a link to the IMSM container device during incremental assembly. This happened when the device metadata did not provide a name for the container. With this update, if no container name has been provided, the metadata version name is used as the container name and a digit is added to the end of the container name so that the link to the IMSM container device is created correctly.

BZ#754998

When an array reshape process was restarted during array assembly, the file system placed on the array could not be mounted and the system returned a busy error. This happened because the child process of the reshape that handled the external metadata of the array was not closed on restart and a memory leak occurred. Consequently, the metadata update failed. With this update, the underlying code has been changed so that the child process of the reshape process is closed under these circumstances and the problem no longer occurs.

BZ#754986

When migrating a non-RAID system to a RAID5 system, an excessive amount of memory could be consumed and the migration process could fail due to a memory leak. With this update, the underlying code has been modified so that the respective resources are freed and the problem no longer occurs.

BZ#812001

If a system containing IMSM RAID devices was rebooted during the rebuild of the IMSM RAID devices, the MD driver changed the device sync_action status from **recover** to **idle**. Consequently, the **mdmon** daemon could detect this change, finish the rebuild process, and write the metadata of the unfinished rebuild process to disks before the restart. After restart, the RAID volume was in the **Normal** state in OROM and the rebuild seemed to be finished. However, the RAID volume was in the **auto-read-only** state, metadata was in the **Dirty** state, and the data was inconsistent (out-of-sync). With this update, the appropriate test has been added and, when mdmon now detects the change of sync_action from **recover** to **idle**, it checks if the rebuild process has really finished.

BZ#814743

An entry could be missing in the device map file if an entry for an array with the same name existed in the device map file. This update fixes the problem so the entry is added to the device map file in such cases.

Enhancement

BZ#808475

The **--size max** option has been added, allowing the administrator of an IMSM array to enlarge the last volume in the array to take up the remaining space in the array.

Users of mdadm should upgrade to these updated packages, which fix these bugs and add these enhancements.

5.189. METACITY

5.189.1. RHBA-2012:0994 – metacity bug fix update

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

Metacity is a window manager that integrates with the GNOME desktop.

BZ#802916

Previously, password dialog boxes did not automatically have the focus after they were opened and could remain covered by a full-screen window. With this update, password dialog boxes are automatically given the focus so that they are raised above full-screen windows, and the problem no longer occurs.

All users of metacity are advised to upgrade to these updated packages, which fix this bug.

5.190. MICROCODE_CTL

5.190.1. RHBA-2012:0827 – microcode_ctl bug fix and enhancement update

Updated microcode_ctl packages that fix one bug and add two enhancements are now available for Red Hat Enterprise Linux 6.

The microcode_ctl packages provide microcode updates for Intel and AMD processors.

Bug Fix

BZ#768803

Previously, running the microcode_ctl utility with long arguments for the "-d" or "-f" options led to a buffer overflow. Consequently, microcode_ctl terminated unexpectedly with a segmentation fault and a backtrace was displayed. With this update, microcode_ctl has been modified to handle this situation gracefully. The microcode_ctl utility no longer crashes and displays an error message informing the user that the file name used is too long.

Enhancements

BZ#736266

The Intel CPU microcode file has been updated to version 20111110, which is the latest version of the microcode available from Intel.

BZ#787757

The AMD CPU microcode file has been updated to version 20120117, which is the latest version of the microcode available from AMD.

All users of microcode_ctl are advised to upgrade to these updated packages, which fix this bug and add these enhancements.

5.191. MINGW32-LIBXML2

5.191.1. RHSA-2013:0217 – Important: mingw32-libxml2 security update

Updated mingw32-libxml2 packages that fix several security issues are now available for Red Hat Enterprise Linux 6. This advisory also contains information about future updates for the mingw32 packages, as well as the deprecation of the packages with the release of Red Hat Enterprise Linux 6.4.

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

These packages provide the libxml2 library, a development toolbox providing the implementation of various XML standards, for users of MinGW (Minimalist GNU for Windows).

Security Fixes

CVE-2011-3919



IMPORTANT

The mingw32 packages in Red Hat Enterprise Linux 6 will no longer be updated proactively and will be deprecated with the release of Red Hat Enterprise Linux 6.4. These packages were provided to support other capabilities in Red Hat Enterprise Linux and were not intended for direct customer use. Customers are advised to not use these packages with immediate effect. Future updates to these packages will be at Red Hat's discretion and these packages may be removed in a future minor release.

A heap-based buffer overflow flaw was found in the way libxml2 decoded entity references with long names. A remote attacker could provide a specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2012-5134

A heap-based buffer underflow flaw was found in the way libxml2 decoded certain entities. A remote attacker could provide a specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2012-0841

It was found that the hashing routine used by libxml2 arrays was susceptible to predictable hash collisions. Sending a specially-crafted message to an XML service could result in longer processing time, which could lead to a denial of service. To mitigate this issue, randomization has been added to the hashing function to reduce the chance of an attacker successfully causing intentional collisions.

CVE-2010-4008, **CVE-2010-4494**, **CVE-2011-2821**, **CVE-2011-2834**

Multiple flaws were found in the way libxml2 parsed certain XPath (XML Path Language) expressions. If an attacker were able to supply a specially-crafted XML file to an application using libxml2, as well as an XPath expression for that application to run against the crafted file, it could cause the application to crash.

CVE-2011-0216, **CVE-2011-3102**

Two heap-based buffer overflow flaws were found in the way libxml2 decoded certain XML files. A remote attacker could provide a specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2011-1944

An integer overflow flaw, leading to a heap-based buffer overflow, was found in the way libxml2 parsed certain XPath expressions. If an attacker were able to supply a specially-crafted XML file to an application using libxml2, as well as an XPath expression for that application to run against the crafted file, it could cause the application to crash or, possibly, execute arbitrary code.

CVE-2011-3905

An out-of-bounds memory read flaw was found in libxml2. A remote attacker could provide a specially-crafted XML file that, when opened in an application linked against libxml2, would cause the application to crash.

Red Hat would like to thank the Google Security Team for reporting the CVE-2010-4008 issue. Upstream acknowledges Bui Quang Minh from Bkis as the original reporter of CVE-2010-4008.

All users of mingw32-libxml2 are advised to upgrade to these updated packages, which contain backported patches to correct these issues.

5.192. MINGW32-MATAHARI

5.192.1. RHBA-2012:0984 – mingw32-matahari bug fix update

An updated mingw32-matahari package that fixes one bug is now available for Red Hat Enterprise Linux 6.

This package includes Matahari Qpid Management Framework (QMF) Agents for Windows guests. QMF Agent can be used to control and manage various pieces of functionality for an ovirt node, using the Advanced Message Queuing Protocol (AMQP) protocol.

Bug Fix

BZ#806948

Previously, Matahari depended on libqpidclient and libqpidcommon. As a result, Qpid's APIs using libqpidclient and libqpidcommon did not have stable ABI and rebuilding Qpid negatively affected mingw32-matahari. With this update, the dependencies have been removed, thus fixing this bug.

All mingw32-matahari users are advised to upgrade to this updated package, which fixes this bug.

5.193. MINGW32-QPID-CPP

5.193.1. RHBA-2012:0756 – mingw32-qpid-cpp bug fix update

Updated mingw32-qpid-cpp packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The mingw32-qpid-cpp packages provide a message broker daemon that receives, stores, and routes messages by means of runtime libraries for Advanced Message Queuing Protocol (AMQP) client applications developed using the Qpid C++ language.

Bug Fixes

BZ#751349

Previously, HTML documentation was required by the mingw32-qpid-cpp package builds, but was not available. Consequently, the following error message was displayed during the build process:

```
CMake Error at docs/api/cmake_install.cmake:31 (FILE): file INSTALL cannot find file
"/usr/src/redhat/BUILD/qpid-cpp-0.12/build/docs/api/html" to install.
```

As the HTML documentation is not considered essential, this update disables its generation. As a result, the aforementioned error message is not displayed during the build process of the mingw32-qpid-cpp package.

BZ#807345

Previously, mingw32-qpid-cpp had an unnecessary dependency on the mingw32-gnutls package. This update removes the dependency.

BZ#813537

Previously, mingw32-qpid-cpp had an unnecessary dependency on the mingw32-libxslt package. This update removes the dependency.

All users of mingw32-qpid-cpp are advised to upgrade to these updated packages, which fix these bugs.

5.194. MKBOOTDISK

5.194.1. RHBA-2012:0403 – mkbootdisk bug fix update

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

The mkbootdisk program creates a standalone boot floppy disk for booting the running system.

Bug Fixes

BZ#761590

Previously, the "mkbootdisk --iso" command could fail with a "Volume ID string too long" error. This was because the length of the volume ID is limited to 32 characters but the limit was not verified. With this update, the string length is verified and if it is longer than 32 characters, the fixed string "Red Hat Linux" is used for volume ID.

BZ#790039

Previously, the mkbootdisk package was missing a dependency on the genisoimage package that was required for functionality of the "--iso" option. This update adds the missing dependency.

All users of mkbootdisk are advised to upgrade to this updated package, which fixes these bugs.

5.195. MLOCATE

5.195.1. RHBA-2012:1355 – mlocate bug fix update

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

The mlocate packages provide a locate/updatedb implementation. Mlocate keeps a database of all existing files and allows you to look up files by name.

Bug Fixes

BZ#690800

Prior to this update, the locate(1) manual page contained a misprint. This update corrects the misprint.

BZ#699363

Prior to this update, the mlocate tool aborted the "updatedb" command if an incorrect filesystem implementation returned a zero-length file name. As a consequence, the locate database was not be updated. This update detects invalid zero-length file names, warns about them, and continues to the locate database.

All users of mlocate are advised to upgrade to these updated packages, which fix these bugs.

5.196. MOD_AUTH_KERB

5.196.1. RHBA-2012:0877 – mod_auth_kerb bug fix and enhancement update

Updated mod_auth_kerb packages that fix a bug and add an enhancement 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#688210

Due to a bug in the handling of memory lifetime when the module was configured to allow delegated credentials, the $KRB5CCNAME variable was lost after the first request of an authenticated connection, causing web applications which relied on the presence of delegated credentials to fail. The memory lifetime handling has been fixed, allowing such web applications to access delegated credentials.

Enhancement

BZ#767741

Support for "S4U2Proxy" constrained delegation has been added, which allows mod_auth_kerb to obtain credentials on behalf of an authenticated user.

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

5.197. MOD_AUTHZ_LDAP

5.197.1. RHBA-2012:1389 – mod_authz_ldap bug fix update

Updated mod_authz_ldap packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The mod_authz_ldap packages provide a module for the Apache HTTP Server to authenticate users against an LDAP database.

Bug Fixes

BZ#607797

Prior to this update, the License field of the mod_authz_ldap packages contained an incorrect tag. This update modifies the license text. Now, the license tag correctly reads "ASL1.0".

BZ#643691

Prior to this update, the mod_authz_ldap module could leak memory. As a consequence, the memory consumption of the httpd process could increase as more requests were processed. This update modifies the underlying code to handle LDAP correctly. Now, the memory consumption as at

expected levels.

BZ#782442

Prior to this update, passwords were logged in plain text to the error log when an LDAP bind password was configured if a connection error occurred. This update modifies the underlying code to prevent passwords from being logged in error conditions.

All users of mod_authz_ldap are advised to upgrade to this updated package, which fixes these bugs.

5.198. MOD_NSS

5.198.1. RHBA-2012:0919 – mod_nss bug fix update

An updated mod_nss package that fixes several bugs is now available for Red Hat Enterprise Linux 6.

The mod_nss module provides strong cryptography for the Apache HTTP Server via the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols, using the Network Security Services (NSS) security library.

Bug Fixes

BZ#749408

PK11_ListCerts was called for every server instead of only once. If there were more than a few hundred certificates in the database, the PK11_ListCerts call could take several seconds or even minutes.

BZ#749409

ECC is now enabled by default for mod_nss.

BZ#797326 and BZ#797358

The fix for BZ#691502, related to clearing the SSL cache when mod_nss started, introduced a file descriptor leak in the httpd Apache daemon. This has been fixed.

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

5.199. MODULE-INIT-TOOLS

5.199.1. RHBA-2012:0871 – module-init-tools bug fix and enhancement update

Updated module-init-tools packages that fix two bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The module-init-tools packages include various programs needed for automatic loading and unloading of modules under 2.6 kernels, as well as other module management programs. Device drivers and file systems are two examples of loaded and unloaded modules.

Bug Fixes

BZ#670613

Previously, on low-memory systems (such as low-memory high-performance infrastructure, or HPC,

nodes or virtual machines), depmod could use excessive amount of memory. As a consequence, the depmod process was killed by the OOM (out of memory) mechanism, and the system was unable to boot. With this update, the free() function is correctly used on several places in the code so that depmod's memory consumption is reduced.

BZ#673100

Previously, if the "override" keyword was present in the depmod.conf file without any parameters specified, the depmod utility terminated unexpectedly with a segmentation fault. A patch has been applied to ensure that the depmod utility no longer crashes and a syntax warning is displayed instead.

Enhancement

BZ#761511

This update adds the "backports" directory to the search path in the depmod.conf file, which is necessary to support integration of the compat-wireless package into kernel packages.

All users of module-init-tools are advised to upgrade to these updated packages, which fix these bugs and add this enhancements.

5.200. MOD_WSGI

5.200.1. RHBA-2012:1358 – mod_wsgi bug fix and enhancement update

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

The mod_wsgi packages provide a Apache httpd module, which implements a WSGI compliant interface for hosting Python based web applications.

Bug Fix

BZ#670577

Prior to this update, a misleading warning message from the mod_wsgi utilities was logged during startup of the Apache httpd daemon. This update removes this message from the mod_wsgi module.

Enhancement

BZ#719409

With this update, access to the SSL connection state is now available in WSGI scripts using the methods "mod_ssl.is_https" and "mod_ssl.var_lookup".

All users of mod_wsgi are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

5.201. MRTG

5.201.1. RHBA-2012:1449 – mrtg bug fix update

Updated mrtg packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The mrtg packages provide the Multi Router Traffic Grapher (MRTG) to monitor the traffic load on network-links. MRTG generates HTML pages containing Portable Network Graphics (PNG) images, which provide a live, visual representation of this traffic.

Bug Fixes

BZ#707188

Prior to this update, the MRTG tool did not handle the socket6 correctly. As a consequence, MRTG reported errors when run on a system with an IPv6 network interface due to a socket conflict. This update modifies the underlying code to socket6 as expected. (#706519)

* Prior to this update, changing the "kMG" keyword in the MRTG configuration could cause the labels on the y-axis to overlap the main area of the generated chart. With this update, an upstream patch has been applied to address this issue, and changing the "kMG" keyword in the configuration no longer leads to the incorrect rendering of the resulting charts.

BZ#836197

Prior to this update, the wrong value was returned from the IBM Fibrechannel switch when using the ifSpeed interface. As a consequence, mrtg cfgmaker failed to use ifHighSpeed on IBM FibreChannel switches. This update modifies the underlying code to return the correct value.

All users of mrtg are advised to upgrade to these updated packages, which fix these bugs.

5.202. MT-ST

5.202.1. RHBA-2012:1409 – mt-st bug fix update

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

The mt-st package contains the mt and st tape drive management programs. Mt (for magnetic tape drives) and st (for SCSI tape devices) can control rewinding, ejecting, skipping files and blocks and more.

Bug Fix

BZ#820245

Prior this update, the stinit init script did not support standard actions like "status" or "restart". As a consequence, an error code was returned. This update modifies the underlying code to use all use all standard actions.

All users of mt-st are advised to upgrade to these updated packages, which fix this bug.

5.203. MYSQL-CONNECTOR-JAVA

5.203.1. RHBA-2012:0992 – mysql-connector-java bug fix and enhancement update

An updated mysql-connector-java package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The mysql-connector-java package provides a native Java driver that converts JDBC (Java Database Connectivity) calls into the network protocol used by the MySQL database. It lets developers working with the Java programming language easily build programs and applets that interact with MySQL.

The mysql-connector-java package has been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version. For a list of changes, refer to the MySQL Connector/J documentation.

This update also adds "stub" implementations of methods required by the JDBC 4.1 API specification. Currently, these methods throw exceptions when called, but their presence is necessary for the driver to function properly in JDK 7 and later. This update also converts the driver from a GCJ build to a pure jar (noarch) build. (BZ#816696)

Users are advised to upgrade to this updated mysql-connector-java package, which resolves these issues and adds these enhancements.

5.204. MYSQL

5.204.1. RHSA-2012:1551 – Important: mysql security update

Updated mysql packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

MySQL is a multi-user, multi-threaded SQL database server. It consists of the MySQL server daemon (mysqld) and many client programs and libraries.

Security Fix

CVE-2012-5611

A stack-based buffer overflow flaw was found in the user permission checking code in MySQL. An authenticated database user could use this flaw to crash the mysqld daemon or, potentially, execute arbitrary code with the privileges of the user running the mysqld daemon.

All MySQL users should upgrade to these updated packages, which correct this issue. After installing this update, the MySQL server daemon (mysqld) will be restarted automatically.

5.204.2. RHSA-2012:1462 – Important: mysql security update

Updated mysql packages that fix several security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

MySQL is a multi-user, multi-threaded SQL database server. It consists of the MySQL server daemon (mysqld) and many client programs and libraries.

Security Fix

CVE-2012-1688, CVE-2012-1690, CVE-2012-1703, CVE-2012-2749, CVE-2012-0540, CVE-2012-1689, CVE-2012-1734, CVE-2012-3163, CVE-2012-3158, CVE-2012-3177, CVE-2012-3166, CVE-2012-3173, CVE-2012-3150, CVE-2012-3180, CVE-2012-3167, CVE-2012-3197, CVE-2012-3160

This update fixes several vulnerabilities in the MySQL database server. Information about these flaws can be found on the Oracle Critical Patch Update Advisory pages:

http://www.oracle.com/technetwork/topics/security/cpuapr2012-366314.html#AppendixMSQLhttp://www.oracle.com/technetwork/topics/security/cpujul2012-392727.html#AppendixMSQLhttp://www.oracle.com/technetwork/topics/security/cpuoct2012-1515893.html#AppendixMSQL

These updated packages upgrade MySQL to version 5.1.66. Refer to the MySQL release notes for a full list of changes:

http://dev.mysql.com/doc/refman/5.1/en/news-5-1-62.htmlhttp://dev.mysql.com/doc/refman/5.1/en/news-5-1-63.htmlhttp://dev.mysql.com/doc/refman/5.1/en/news-5-1-64.htmlhttp://dev.mysql.com/doc/refman/5.1/en/news-5-1-65.htmlhttp://dev.mysql.com/doc/refman/5.1/en/news-5-1-66.html

All MySQL users should upgrade to these updated packages, which correct these issues. After installing this update, the MySQL server daemon (mysqld) will be restarted automatically.

5.204.3. RHSA-2013:0219 – Moderate: mysql security update

Updated mysql packages that fix several security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

MySQL is a multi-user, multi-threaded SQL database server. It consists of the MySQL server daemon (mysqld) and many client programs and libraries.

Security Fix

CVE-2012-0572, CVE-2012-0574, CVE-2012-1702, CVE-2012-1705, CVE-2013-0375, CVE-2013-0383, CVE-2013-0384, CVE-2013-0385, CVE-2013-0389

This update fixes several vulnerabilities in the MySQL database server. Information about these flaws can be found on the Oracle Critical Patch Update Advisory page:

http://www.oracle.com/technetwork/topics/security/cpujan2013-1515902.html#AppendixMSQL

These updated packages upgrade MySQL to version 5.1.67. Refer to the MySQL release notes for a full list of changes:

http://dev.mysql.com/doc/relnotes/mysql/5.1/en/news-5-1-67.html

All MySQL users should upgrade to these updated packages, which correct these issues. After installing this update, the MySQL server daemon (mysqld) will be restarted automatically.

5.204.4. RHSA-2012:0874 – Low: mysql security and enhancement update

Updated mysql packages that fix one security issue 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 low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

MySQL is a multi-user, multi-threaded SQL database server. It consists of the MySQL server daemon (mysqld) and many client programs and libraries.

Security Fix

CVE-2012-2102

A flaw was found in the way MySQL processed HANDLER READ NEXT statements after deleting a record. A remote, authenticated attacker could use this flaw to provide such requests, causing mysqld to crash. This issue only caused a temporary denial of service, as mysqld was automatically restarted after the crash.

Enhancement

BZ#740224

The InnoDB storage engine is built-in for all architectures. This update adds InnoDB Plugin, the InnoDB storage engine as a plug-in for the 32-bit x86, AMD64, and Intel 64 architectures. The plug-in offers additional features and better performance than when using the built-in InnoDB storage engine. Refer to the MySQL documentation, for information about enabling the plug-in.

All MySQL users should upgrade to these updated packages, which add this enhancement and contain a backported patch to correct this issue. After installing this update, the MySQL server daemon (mysqld) will be restarted automatically.

5.205. NAUTILUS

5.205.1. RHBA-2012:0914 – nautilus bug fix update

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

Nautilus is the file manager and graphical shell for the GNOME desktop that makes it easy to manage your files and the rest of your system. It allows to browse directories on local and remote file systems, preview files and launch applications associated with them. It is also responsible for handling the icons on the GNOME desktop.

Bug Fixes

BZ#600260

When the display size was changed, no desktop folder refresh was performed. Consequently, icons disappeared outside the visible screen when the screen resolution was lowered. With this update, an explicit refresh action has been placed on screen size changes and the icons are now visible when the screen resolution is lowered.

BZ#782467

Previously, an empty file with the name ending ".desktop" was automatically identified as a special file. Consequently, some operations on the file, such as rename, failed. Now, a fallback that allows to

use the regular file rename operation has been added to the code and these files can now be renamed as expected.

BZ#772103

Due to a short-lived internal object, free-space information was not displayed in the volume Properties dialog. With this update, a reference to another internal object has been placed in the code and the free-space information is now displayed properly.

BZ#755561

Previously, when the Nautilus desktop was set to display a user home directory, an internal queued load operation did not get canceled after refresh. Consequently, nautilus terminated unexpectedly on startup. With this update, pending internal operations that are not valid are correctly canceled and the crashes no longer occur.

All users of nautilus are advised to upgrade to these updated packages, which fix these bugs.

5.206. NET-SNMP

5.206.1. RHBA-2012:1106 – net-snmp bug fix update

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

The net-snmp packages provide various libraries and tools for the Simple Network Management Protocol (SNMP), including an SNMP library, an extensible agent, tools for requesting or setting information from SNMP agents, tools for generating and handling SNMP traps, a version of the netstat command which uses SNMP, and a Tk/Perl Management Information Base (MIB) browser.

Bug Fix

BZ#836252

Prior to this update, there was a limit of 50 'exec' entries in the /etc/snmp/snmpd.conf file. With more than 50 such entries in the configuration file, the snmpd daemon returned the "Error: No further UCD-compatible entries" error message to the system log. With this update, this limit has been removed and there can now be any number of 'exec' entries in the snmpd configuration file, thus preventing this bug.

All users of net-snmp are advised to upgrade to these updated packages, which fix this bug.

5.206.2. RHSA-2012:0876 – Moderate: net-snmp security and bug fix update

Updated net-snmp packages that fix one security issue and multiple 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

The net-snmp packages provide various libraries and tools for the *Simple Network Management Protocol* (SNMP), including an **SNMP** library, an extensible agent, tools for requesting or setting information from **SNMP** agents, tools for generating and handling SNMP traps, a version of the **netstat** command which uses SNMP, and a **Tk/Perl** *Management Information Base* (MIB) browser.

Security Fix

CVE-2012-2141

An array index error, leading to an out-of-bounds buffer read flaw, was found in the way the **net-snmp** agent looked up entries in the extension table. A remote attacker with read privileges to a Management Information Base (MIB) subtree handled by the ***extend*** directive (in **/etc/snmp/snmpd.conf**) could use this flaw to crash **snmpd** via a crafted SNMP GET request.

Bug Fixes

BZ#736580

In the previous update, a change was made in order to stop **snmpd** terminating unexpectedly when an **AgentX** subagent disconnected while processing a request. This fix, however, introduced a memory leak. With this update, this memory leak is fixed.

BZ#740172

In a previous update, a new BRIDGE-MIB was implemented in the net-snmp-perl subpackage. This MIB used incorrect conversion of interface-index values from the kernel and reported incorrect values of ifIndex ***OIDs*** (object identifiers). With this update, conversion of interface indexes is fixed and BRIDGE-MIB reports correct ifIndex OIDs.

BZ#746903

Previously, **snmpd** erroneously enabled verbose logging when parsing the **proxy** option in the **snmpd.conf** file. Consequently, unexpected debug messages were sometimes written to the system log. With this update, **snmpd** no longer modifies logging settings when parsing the **proxy** option. As a result, no debug messages are sent to the system log unless explicitly enabled by the system administrator.

BZ#748410

Previously, the **snmpd** daemon strictly implemented *RFC 2780*. However, this specification no longer scales well with modern big storage devices with small allocation units. Consequently, **snmpd** reported a wrong value for the "HOST-RESOURCES-MIB::hrStorageSize" object when working with a large file system (larger than 16TB), because the accurate value did not fit into Integer32 as specified in the RFC. To address this problem, this update adds a new option to the **/etc/snmp/snmpd.conf** configuration file, "realStorageUnits". By changing the value of this option to **0**, users can now enable recalculation of all values in "hrStorageTable" to ensure that the multiplication of "hrStorageSize" and "hrStorageAllocationUnits" always produces an accurate device size. The values of "hrStorageAllocationUnits" are then artificial in this case and no longer represent the real size of the allocation unit on the storage device.

BZ#748411, BZ#755481, BZ#757685

In the previous net-snmp update, the implementation of "HOST-RESOURCES-MIB::hrStorageTable" was rewritten and devices with *Veritas File System* (VxFS), *ReiserFS*, and *Oracle Cluster File System* (OCFS2) were not reported. In this update, **snmpd** properly recognizes VxFS, ReiserFS, and OCFS2 devices and reports them in "HOST-RESOURCES-MIB::hrStorageTable".

BZ#748907

Prior to this update, the Net-SNMP Perl module did not properly evaluate error codes in the **register()** method in the "NetSNMP::agent" module and terminated unexpectedly when this method failed. With this update, the **register()** method has been fixed and the updated Perl modules no

longer crash on failure.

BZ#749227

The SNMP daemon (**snmpd**) did not properly fill a set of watched socket file descriptors. Therefore, the daemon sometimes terminated unexpectedly with the "select: bad file descriptor" error message when more than 32 **AgentX** subagents connected to **snmpd** on 32-bit platforms or more than 64 subagents on 64-bit platforms. With this update, **snmpd** properly clears sets of watched file descriptors and no longer crashes when handling a large number of subagents.

BZ#754275

Previously, **snmpd** erroneously checked the length of "SNMP-TARGET-MIB::snmpTargetAddrRowStatus" value in incoming "SNMP-SET" requests on 64-bit platforms. Consequently, **snmpd** sent an incorrect reply to the "SNMP-SET" request. With this update, the check of "SNMP-TARGET-MIB::snmpTargetAddrRowStatus" is fixed and it is possible to set it remotely using "SNMP-SET" messages.

BZ#754971

Previously, **snmpd** did not check the permissions of its MIB index files stored in the **/var/lib/net-snmp/mib_indexes** directory and assumed it could read them. If the read access was denied, for example due to incorrect SELinux contexts on these files, **snmpd** crashed. With this update, **snmpd** checks if its MIB index files were correctly opened and does not crash if they cannot be opened.

BZ#786931

Before this release, the length of the ***OID*** parameter of "sysObjectID" (an **snmpd.conf** config file option) was not correctly stored in **snmpd**, which resulted in "SNMPv2-MIB::sysObjectID" being truncated if the ***OID*** had more than 10 components. In this update, handling of the ***OID*** length is fixed and "SNMPv2-MIB::sysObjectID" is returned correctly.

BZ#788954

Prior to this update, when **snmpd** was started and did not find a network interface which had been present during the last **snmpd** shutdown, the following error message was logged:

```
snmpd: error finding row index in _ifXTable_container_row_restore
```

This happened on systems which dynamically create and remove network interfaces on demand, such as virtual hosts or **PPP** servers. In this update, this message has been removed and no longer appears in the system log.

BZ#789909

Previously, **snmpd**, enumerated active **TCP** connections for "TCP-MIB::tcpConnectionTable" in an inefficient way with O(n^2) complexity. With many TCP connections, an **SNMP** client could time out before **snmpd** processed a request regarding the "tcpConnectionTable", and sent a response. This update improves the enumeration mechanism and **snmpd** now swiftly responds to SNMP requests in the "tcpConnectionTable".

BZ#799291

When an object identifier (***OID***) was out of the subtree registered by the proxy statement in the **/etc/snmp/snmpd.conf** configuration file, the previous version of the **snmpd** daemon failed to use a correct ***OID*** of proxied "GETNEXT" requests. With this update, snmpd now adjusts the ***OIDs*** of proxied "GETNEXT" requests correctly and sends correct requests to the remote agent as expected.

BZ#822480

Net-SNMP daemons and utilities use the **/var/lib/net-snmp** directory to store persistent data, for example the cache of parsed MIB files. This directory is created by the net-snmp package and when this package is not installed, Net-SNMP utilities and libraries create the directory with the wrong SELinux context, which results in an Access Vector Cache (AVC) error reported by SELinux. In this update, the **/var/lib/net-snmp** directory is created by the net-snmp-lib package, therefore all Net-SNMP utilities and libraries do not need to create the directory and the directory will have the correct SELinux context.

All users of net-snmp are advised to upgrade to these updated packages, which contain backported patches to resolve these issues. After installing the update, the **snmpd** and **snmptrapd** daemons will be restarted automatically.

5.206.3. RHBA-2013:1111 – net-snmp bug fix update

Updated net-snmp packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The net-snmp packages provide various libraries and tools for the Simple Network Management Protocol (SNMP), including an SNMP library, an extensible agent, tools for requesting or setting information from SNMP agents, tools for generating and handling SNMP traps, a version of the netstat command which uses SNMP, and a Tk/Perl Management Information Base (MIB) browser.

Bug Fix

BZ#986192

In previous Net-SNMP releases, snmpd reported an invalid speed of network interfaces in IF-MIB::ifTable and IF-MIB::ifXTable if the interface had a speed other than 10, 100, 1000 or 2500 MB/s. Thus, the net-snmp ifHighSpeed value returned 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 ethtool, and snmpd correctly reports non-standard network interface speeds.

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

5.206.4. RHBA-2013:1216 – net-snmp bug fix update

Updated net-snmp packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The net-snmp packages provide various libraries and tools for the Simple Network Management Protocol (SNMP), including an SNMP library, an extensible agent, tools for requesting or setting information from SNMP agents, tools for generating and handling SNMP traps, a version of the netstat command which uses SNMP, and a Tk/Perl Management Information Base (MIB) browser.

Bug Fix

BZ#1002859

When an AgentX subagent disconnected from the SNMP daemon (snmpd), the daemon did not properly check that there were no active requests queued in the subagent and destroyed the session. Consequently, the session was referenced by snmpd later when processing queued requests

and because it was already destroyed, snmpd terminated unexpectedly with a segmentation fault or looped indefinitely. This update adds several checks to prevent the destruction of sessions with active requests, and snmpd no longer crashes in the described scenario.

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

5.207. NETWORKMANAGER-OPENSWAN

5.207.1. RHBA-2012:0915 – NetworkManager-openswan bug fix update

An updated NetworkManager-openswan package that fixes various bugs is now available for Red Hat Enterprise Linux 6.

NetworkManager-openswan contains software for integrating the Openswan VPN software with NetworkManager and the GNOME desktop.

Bug Fixes

BZ#696946

Prior to this update, it was possible to enter an incorrect IP address into the gateway field and it was sometimes interpreted as a gateway hostname. With this update, the code has been improved to validate IP addresses. As a result, valid IPv4 addresses, as well as hostnames, are more reliably distinguished.

BZ#748365

NetworkManager-openswan was not able to import configuration files that were previously exported using NetworkManager. This release adds support for this functionality and importing Openswan IPsec configuration files is now possible using NetworkManager.

All users of NetworkManager-openswan are advised to upgrade to this updated package, which fixes these bugs.

5.208. NETWORKMANAGER

5.208.1. RHBA-2012:1518 – NetworkManager bug fix and enhancement update

Updated NetworkManager packages that fix three bugs and add an enhancement are now available for Red Hat Enterprise Linux 6.

NetworkManager is a system network service that manages network devices and connections, attempting to keep active network connectivity when available. It manages Ethernet, wireless, mobile broadband (WWAN), and PPPoE (Point-to-Point Protocol over Ethernet) devices, and provides VPN integration with a variety of different VPN services.

Bug Fixes

BZ#864828

Due to a bug reading and writing network configuration files, network connections using the LEAP authentication method could not be made available to all users. A patch has been provided to address this issue and the network configuration files now allow LEAP as expected.

BZ#864829

Previously, NetworkManager did not allow selecting the WPA protocol version for connection. Certain enterprise WLAN networks using Cisco equipment do not allow roaming between WPA and WPA2 Virtual Access Points (VAP) provided by the same physical access point, requiring the use of a specific WPA protocol version to prevent disconnections. This update adds a WPA protocol combo box to the NetworkManager user interface (UI), thus allowing a specific WPA protocol version to be used when necessary and preventing this bug.

BZ#864830

When a connection was locked to a specific WPA protocol version (either v1 or v2/RSN) via either the GConf system or settings in the /etc/sysconfig/network-scripts/ configuration files, the NetworkManager UI overwrote that preference when the connection was edited and saved. This bug has been fixed and such WPA preferences are now preserved in the described scenario.

Enhancement

BZ#864831

Support for Opportunistic Key Caching (OKC), also known as Proactive Key Caching (PKC), has been added to NetworkManager for all WPA-Enterprise configurations.

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

5.208.2. RHBA-2012:0832 – NetworkManager bug fix and enhancement update

An updated NetworkManager package that fixes several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

NetworkManager is a system network service that manages network devices and connections, attempting to keep active network connectivity when available. It manages Ethernet, wireless, mobile broadband (**WWAN**), and **PPPoE** devices, and provides **VPN** integration with a variety of different VPN services.

Bug Fixes

BZ#663820

NetworkManager used a **DHCP** transaction timeout of 45 seconds without the possibility of configuring a different value. Consequently, in certain cases NetworkManager failed to obtain a network address. NetworkManager has been extended to read the timeout parameter from a DHCP configuration file and use that instead of the default value. As a result, NetworkManager will wait to obtain an address for the duration of the DHCP transaction period as specified in the DHCP client configuration file.

BZ#696967

If no Web browser was installed and the web site link was clicked in **nm-applet**'s **About** dialog box, there was no response and **NetworkManager** did not display an error message. This has now been improved and the user is now presented with an error dialog box in the scenario described.

BZ#747649

NetworkManager did not update the timestamps for system connections. Consequently, the entry under the heading "Last Used" in the connection editor was always "never", even if the connection

had been used. An upstream patch has been applied and NetworkManager now updates connection timestamps for all connections.

BZ#773590

After system boot or when the user re-logged into the Gnome display manager, **NetworkManager** tried to initialize the user settings proxy even if it was not active. Consequently, this caused an unnecessary warning to be written to the **/var/log/messages** file. An upstream patch has been applied to prevent NetworkManager from trying to initialize the user settings proxy if the user settings service does not exist. As a result, warning messages are no longer generated in the scenario described.

BZ#787084

NetworkManager inserted erroneous warning messages in the **/var/log/messages** log file when changing the hostname. An upstream patch has been applied to the **nm-dispatcher** script and NetworkManager no longer generates unnecessary warnings during a hostname change.

BZ#801744

When an existing DHCP lease was renewed, NetworkManager did not recognize it as a change in DHCP state and therefore failed to run the dispatcher scripts. Consequently, hostnames where purged from DHCP records. With this update the code has been improved and NetworkManager now handles same-state transitions correctly. As a result, hostnames are not purged from the DHCP server when a lease is renewed.

BZ#809784

There were errors in **nm-applet**'s message catalog for some languages. Consequently, the **Routes** button name in the connection editor was not translated for those languages and appeared in English. The message catalog has been corrected and now the button text is translated correctly for all supported languages.

Enhancements

BZ#209339

NetworkManager did not support **EAP-FAST** authentication for **WPA2 Enterprise** wireless networks, which made it unusable in some wireless environments. NetworkManager has been enhanced to handle EAP-FAST authentication.

BZ#673476

NetworkManager did not handle *RFC3442*-standard classless static routes provided by a DHCP server without manual changes to the **dhclient**'s configuration file. An enhancement has been made to ensure that RFC3442 classless static routes are requested from the DHCP server, and that they are properly processed by NetworkManager without any manual intervention.

BZ#685096

NetworkManager did not recognize **IP-over-InfiniBand** interfaces, which prevented installation in some situations. These interfaces are now recognized.

BZ#712302, BZ#717475, BZ#804797

Previously, **VLAN** and bonding interfaces were not supported by **NetworkManager** and required special configuration to ensure NetworkManager did not interfere with their operation. NetworkManager now recognizes and can configure VLAN and bonding interfaces but only if the

NM_BOND_VLAN_ENABLED key is set to **yes** in **/etc/sysconfig/network**. The default is that this option is set to **no**.

BZ#719892

If **PolicyKit** setup was used to disable creation of shared **Wi-Fi** networks, and a user tried to create a network using **nm-applet**, the setup silently failed. With this update, nm-applet now issues a notification providing the reason for the failure.

BZ#798294

When no VPN plug-in for **NetworkManager** was installed and a user tried to configure a **VPN** connection, the connection editor displayed an insensitive **Add** button without giving any indication of the cause. This update adds a tooltip to the button informing the user that editing a VPN connection is disabled due to missing VPN plug-ins.

Users are advised to upgrade to these updated NetworkManager packages, which resolve these bugs and add these enhancements.

5.209. NFS4-ACL-TOOLS

5.209.1. RHBA-2012:0783 – nfs4-acl-tools bug fix update

An updated nfs4-acl-tools package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The nfs4-acl-tools package provides utilities for managing NFSv4 Access Control Lists (ACLs) on files and directories mounted on ACL-enabled NFSv4 file systems.

Bug Fix

BZ#769862

Prior to this update, several uninitialized stack pointers were incorrectly freed. As a consequence, the "nfs4_setfacl" command failed with the error message "*** glibc detected *** nfs4_setfacl: double free or corruption (out)" if the format of the input ACL file was incorrect. This update corrects the memory handling process. Now, the nfs4_setfacl command displays a useful error message if the input file syntax is invalid.

All users of nfs4-acl-tools are advised to upgrade to this updated package, which fixes this bug.

5.210. NFS-UTILS

5.210.1. RHBA-2012:0964 – nfs-utils bug fix update

Updated nfs-utils packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

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#737990

Prior to this update, the nfs(5) man page contained incorrect information on the Transmission Control Protocol (TCP) retries. This update modifies the man page and describes more accurately how the TCP time out code works.

BZ#740472

Prior to this update, the "nfs_rewrite_pmap_mount_options()" function did not interrupt RPC timeouts as expected. As a consequence, mounts that used the "-o bg" and "vers=" options did not retry but failed when the server was down. This update modifies the underlying code to allow mounts to retry when the server is down.

BZ#751089

Prior to this update, the rpc.idmapd daemon handled the "SIGUSR*" signal incorrectly. As a consequence, idmapd could, under certain circumstances, close without an error message. This update modifies the underlying code to process the "SIGUSR*" signal as expected.

BZ#758000

Prior to this update, mount points could not be unmounted when the path contained multiple slash characters. This update modifies the "umount" paths so that the mount point can now be unmounted as expected.

BZ#772543

Prior to this update, nfs-utils used the wrong nfs lock file. As a consequence, the "status nfsd" command did not return the correct status. This update modifies the startup script to use the "/var/lock/subsys/nfsd" file as the nfs lock file. Now the correct nfsd status is returned.

BZ#772619

Prior to this update, NFS ID Mapping could redirect Unicode characters using umlaut diacritics (ö, ä, ü) in group names to the group "nobody". This update deactivates the Unicode characters check.

BZ#787970

Prior to this update, the name mapping daemon idmapd failed to decode group names that contained spaces. This update modifies the character size check for decoding the octal encoded value. Now, group names with spaces are decoded as expected.

BZ#800335

Prior to this update, concurrent executions of the "exportfs" command could, under certain circumstances, cause conflicts when updating the etab file. As a consequence, not all exports were successful. This update modifies the exportfs script to allow for concurrent executions.

BZ#801085

Prior to this update, symlinks mounted with NFS could not be unmounted. This update modifies the underlying code so that symlinks are now exported as expected.

BZ#803946

Prior to this update, the nfsd daemon was started before the mountd daemon and nfsd could not validate file handles with mountd. The NFS client received an "ESTALE" error and client applications failed if an existing client sent requests to the NFS server when nfsd was started. This update changes the startup order of the daemons so that nfsd can use the mountd daemon.

BZ#816149

Prior to this update, the preinstall scriptlet could fail to change the default group ID for nfsnobody. This update modifies the preinstall scriptlet and the default group ID is changed after the nfs-utils upgrade as expected.

BZ#816162

Prior to this update, mounting a subdirectory of non-user accounts could, under certain circumstances, fail. This update modifies the underlying code ensure that also the parent directory of the pseudo exports have root squashing disabled. Now, subdirectories of non-user accounts can be successfully mounted.

All users of nfs-utils are advised to upgrade to these updated packages, which fix these bugs.

5.211. NMAP

5.211.1. RHBA-2012:0817 – nmap bug fix and enhancement update

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

The nmap packages provide a network exploration utility and a security scanner.

The nmap package has been upgraded to version 5.51, which provides a number of bug fixes and enhancements over the previous version and improves performance. (BZ#512042)

Bug Fix

BZ#813734

Prior to this update, the nping man page listed "--md" for the "More Fragments" option. This update corrects this misprint and now correctly lists "--mf" for this option.

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

5.212. NSPLUGINWRAPPER

5.212.1. RHSA-2012:1459 – Low: nspluginwrapper security and bug fix update

Updated nspluginwrapper packages that fix one security issue and one bug 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 for each vulnerability from the CVE link(s) associated with each description below.

nspluginwrapper is a utility which allows 32-bit plug-ins to run in a 64-bit browser environment (a common example is Adobe's browser plug-in for presenting proprietary Flash files embedded in web pages). It includes the plug-in viewer and a tool for managing plug-in installations and updates.

Security Fix

CVE-2011-2486

It was not possible for plug-ins wrapped by nspluginwrapper to discover whether the browser was running in Private Browsing mode. This flaw could lead to plug-ins wrapped by nspluginwrapper using normal mode while they were expected to run in Private Browsing mode.

Bug Fix

BZ#869554

When using the Adobe Reader web browser plug-in provided by the acroread-plugin package on a 64-bit system, opening Portable Document Format (PDF) files in Firefox could cause the plug-in to crash and a black window to be displayed where the PDF should be. Firefox had to be restarted to resolve the issue. This update implements a workaround in nspluginwrapper to automatically handle the plug-in crash, so that users no longer have to keep restarting Firefox.

All users of nspluginwrapper are advised to upgrade to these updated packages, which upgrade nspluginwrapper to upstream version 1.4.4, and correct these issues. After installing the update, Firefox must be restarted for the changes to take effect.

5.213. NSS, NSS-UTIL, AND NSPR

5.213.1. RHSA-2012:0973 – Moderate: nss, nss-util, and nspr security, bug fix, and enhancement update

Updated nss, nss-util, and nspr packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

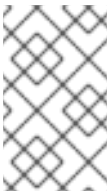
The Red Hat Security Response Team has rated this update as having moderate security impact.

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.

Security Fix

BZ#798533

It was found that a Certificate Authority (CA) issued a subordinate CA certificate to its customer, that could be used to issue certificates for any name. This update renders the subordinate CA certificate as untrusted.



NOTE

This fix only applies to applications using the NSS Builtin Object Token. It does not render the certificates untrusted for applications that use the NSS library, but do not use the NSS Builtin Object Token.

The nspr package has been upgraded to upstream version 4.9, which provides a number of bug fixes and enhancements over the previous version. (BZ#799193)

The nss-util package has been upgraded to upstream version 3.13.3, which provides a number of bug fixes and enhancements over the previous version. (BZ#799192)

The nss package has been upgraded to upstream version 3.13.3, which provides numerous bug fixes and enhancements over the previous version. In particular, SSL 2.0 is now disabled by default, support for SHA-224 has been added, PORT_ErrorToString and PORT_ErrorToName now return the error message and symbolic name of an NSS error code, and NSS_GetVersion now returns the NSS version string. (BZ#744070)

These updated nss, nss-util, and nspr packages also provide fixes for the following bugs:

BZ#746632

A PEM module internal function did not clean up memory when detecting a non-existent file name. Consequently, memory leaks in client code occurred. The code has been improved to deallocate such temporary objects and as a result the reported memory leakage is gone.

BZ#761086

Recent changes to NSS re-introduced a problem where applications could not use multiple SSL client certificates in the same process. Therefore, any attempt to run commands that worked with multiple SSL client certificates, such as the "yum repolist" command, resulted in a re-negotiation handshake failure. With this update, a revised patch correcting this problem has been applied to NSS, and using multiple SSL client certificates in the same process is now possible again.

BZ#768669

The PEM module did not fully initialize newly constructed objects with function pointers set to NULL. Consequently, a segmentation violation in libcurl was sometimes experienced while accessing a package repository. With this update, the code has been changed to fully initialize newly allocated objects. As a result, updates can now be installed without problems.

BZ#784674

A lack-of-robustness flaw caused the administration server for Red Hat Directory Server to terminate unexpectedly because the mod_nss module made nss calls before initializing nss as per the documented API. With this update, nss protects itself against being called before it has been properly initialized by the caller.

BZ#795693

Compilation errors occurred with some compilers when compiling code against NSS 3.13.1. The following error message was displayed:

```
pkcs11n.h:365:26: warning: "__GNUC_MINOR" is not defined
```

An upstream patch has been applied to improve the code and the problem no longer occurs.

BZ#797426

Unexpected terminations were reported in the messaging daemon (qpidd) included in Red Hat Enterprise MRG after a recent update to nss. This occurred because qpidd made nss calls before initializing nss. These updated packages prevent qpidd and other affected processes that call nss without initializing as mandated by the API from crashing.

Users of NSS, NSPR, and nss-util are advised to upgrade to these updated packages, which fix these issues and add these enhancements. After installing this update, applications using NSS, NSPR, or nss-util must be restarted for this update to take effect.

5.213.2. RHSA-2012:1091 – Moderate: nss, nspr, and nss-util security, bug fix, and enhancement update

Updated nss, nss-util, and nspr 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

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.

Security Fix

CVE-2012-0441

A flaw was found in the way the ASN.1 (Abstract Syntax Notation One) decoder in NSS handled zero length items. This flaw could cause the decoder to incorrectly skip or replace certain items with a default value, or could cause an application to crash if, for example, it received a specially-crafted OCSP (Online Certificate Status Protocol) response.



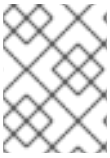
NOTE

The nspr package has been upgraded to upstream version 4.9.1, which provides a number of bug fixes and enhancements over the previous version. (BZ#833762)



NOTE

The nss-util package has been upgraded to upstream version 3.13.5, which provides a number of bug fixes and enhancements over the previous version. (BZ#833763)



NOTE

The nss package has been upgraded to upstream version 3.13.5, which provides a number of bug fixes and enhancements over the previous version. (BZ#834100)

All NSS, NSPR, and nss-util users are advised to upgrade to these updated packages, which correct these issues and add these enhancements. After installing this update, applications using NSS, NSPR, or nss-util must be restarted for this update to take effect.

5.213.3. RHSA-2013:0213 – Important: nss, nss-util, and nspr security, bug fix, and enhancement update

Updated nss, nss-util, and nspr packages that fix one security issue, various bugs, and add enhancements are now available for Red Hat Enterprise Linux 6.

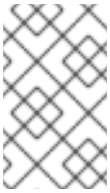
The Red Hat Security Response Team has rated this update as having important security impact.

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.

Security Fix

BZ#890605

It was found that a Certificate Authority (CA) mis-issued two intermediate certificates to customers. These certificates could be used to launch man-in-the-middle attacks. This update renders those certificates as untrusted. This covers all uses of the certificates, including SSL, S/MIME, and code signing.



NOTE

This fix only applies to applications using the NSS Builtin Object Token. It does not render the certificates untrusted for applications that use the NSS library, but do not use the NSS Builtin Object Token.



NOTE

In addition, the nss package has been upgraded to upstream version 3.13.6, the nss-util package has been upgraded to upstream version 3.13.6, and the nspr package has been upgraded to upstream version 4.9.2. These updates provide a number of bug fixes and enhancements over the previous versions. (BZ#891663, BZ#891670, BZ#891661)

Users of NSS, NSPR, and nss-util are advised to upgrade to these updated packages, which fix these issues and add these enhancements. After installing this update, applications using NSS, NSPR, or nss-util must be restarted for this update to take effect.

5.214. NSS-PAM-LDAPD

5.214.1. RHBA-2012:1487 – nss-pam-ldapd bug fix update

Updated nss-pam-ldapd packages that fix a bug is now available for Red Hat Enterprise Linux 6.

The nss-pam-ldapd provides the nss-pam-ldapd daemon (nslcd) which uses a directory server to look up name service information on behalf of a lightweight nsswitch module.

Bug Fix

BZ#864365

When the nslcd daemon requested access to a large group, a buffer provided by the glibc library could not contain such a group and retried again with a larger buffer to process the operation successfully. However, confusing and redundant error messages were written to the /var/log/message file. This update makes sure that even when glibc provides a buffer that is too small on first attempt in the described scenario, no redundant error messages are returned.

All users of nss-pam-ldapd are advised to upgrade to this updated package, which fixes this bug.

5.215. NSS

5.215.1. RHBA-2012:1003 – nss bug fix update

Updated nss packages that fix one bug 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.

Bug Fix

BZ#828679

Due to a missing out-of-memory (OOM) check and improper freeing of allocated memory, the Privacy Enhanced Mail (PEM) module did not fully validate the encoding of certificates stored in a PEM-formatted file. As a consequence, error handling tests failed. With this update, the PEM module correctly validates the encoding, handles memory deallocation consistently, and error handling tests pass as expected.

All users of nss are advised to upgrade to these updated packages, which fix this bug.

5.216. NUMACTL

5.216.1. RHBA-2012:0828 – numactl bug fix and enhancement update

Updated numactl packages that fix multiple bugs and add various enhancements 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.

The numactl packages have been upgraded to upstream version 2.0.7, which provides a number of bug fixes and enhancements over the previous version. (BZ#645066)

Bug Fix

BZ#751764

Prior to this update, the man page for the numastat tool was not included in the numactl package. This update adds the missing numastat man page to numactl.

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

5.217. NUMPY

5.217.1. RHBA-2012:0986 – numpy bug fix and enhancement update

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

The numpy packages provide NumPY. NumPY is an extension to the Python programming language, which adds support for large, multi-dimensional arrays and matrices, and a library of mathematical functions that operate on such arrays.

The numpy packages have been upgraded to upstream version 1.4.1, which provides a number of bug fixes and enhancements over the previous version.

BZ#692959

This update introduces two important changes in NumPY's behavior:

- When operating on 0-d arrays, the numpy.max() function and other functions now no longer accept axis values other than 0, -1, and None, and NumPY now raises an error for other axis values.
- It is now no longer possible to specify an axis value greater than the MAX_DIMS value and NumPY now raises an error under these circumstances.

Refer to the /usr/share/doc/numpy-1.4.1/1.4.0-notes.rst file for further details about the changes.

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

5.218. OPENJPEG

5.218.1. RHSA-2012:1068 – Important: openjpeg security update

Updated openjpeg packages that fix two security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

OpenJPEG is an open source library for reading and writing image files in JPEG 2000 format.

Security Fixes

CVE-2012-3358

An input validation flaw, leading to a heap-based buffer overflow, was found in the way OpenJPEG handled the tile number and size in an image tile header. A remote attacker could provide a specially-crafted image file that, when decoded using an application linked against OpenJPEG, would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

CVE-2009-5030

OpenJPEG allocated insufficient memory when encoding JPEG 2000 files from input images that have certain color depths. A remote attacker could provide a specially-crafted image file that, when opened in an application linked against OpenJPEG (such as image_to_j2k), would cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

Users of OpenJPEG should upgrade to these updated packages, which contain patches to correct these issues. All running applications using OpenJPEG must be restarted for the update to take effect.

5.218.2. RHSA-2012:1283 – Important: openjpeg security update

Updated openjpeg packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

OpenJPEG is an open source library for reading and writing image files in JPEG 2000 format.

Security Fix

CVE-2012-3535

It was found that OpenJPEG failed to sanity-check an image header field before using it. A remote attacker could provide a specially-crafted image file that could cause an application linked against OpenJPEG to crash or, possibly, execute arbitrary code.

This issue was discovered by Huzaifa Sidhpurwala of the Red Hat Security Response Team.

Users of OpenJPEG should upgrade to these updated packages, which contain a patch to correct this issue. All running applications using OpenJPEG must be restarted for the update to take effect.

5.219. OPENLDAP

5.219.1. RHSA-2012:1151 – Low: openldap security and bug fix update

Updated openldap packages that fix one security issue and one bug 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 for each vulnerability from the CVE link(s) associated with each description below.

OpenLDAP is an open source suite of LDAP (Lightweight Directory Access Protocol) applications and development tools.

Security Fix

CVE-2012-2668

It was found that the OpenLDAP server daemon ignored olcTLSCipherSuite settings. This resulted in the default cipher suite always being used, which could lead to weaker than expected ciphers being accepted during Transport Layer Security (TLS) negotiation with OpenLDAP clients.

Bug Fix

BZ#844428

When the smbk5pwd overlay was enabled in an OpenLDAP server, and a user changed their password, the Microsoft NT LAN Manager (NTLM) and Microsoft LAN Manager (LM) hashes were not computed correctly. This led to the sambaLMPassword and sambaNTPassword attributes being updated with incorrect values, preventing the user logging in using a Windows-based client or a Samba client.

With this update, the smbk5pwd overlay is linked against OpenSSL. As such, the NTLM and LM hashes are computed correctly, and password changes work as expected when using smbk5pwd.

Users of OpenLDAP are advised to upgrade to these updated packages, which contain backported patches to correct these issues. After installing this update, the OpenLDAP daemons will be restarted automatically.

5.219.2. RHSA-2012:0899 – Low: openldap bug fix update

Updated openldap packages that fix a 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.

OpenLDAP is an open-source suite of LDAP (Lightweight Directory Access Protocol) applications and development tools. LDAP is a set of protocols for accessing directory services (usually phone-book style information, but other information is possible) over the Internet, similar to the way DNS (Domain Name System) information is propagated over the Internet. The openldap package contains configuration files, libraries, and documentation for OpenLDAP.

Security Fix

CVE-2012-1164

A denial of service flaw was found in the way the OpenLDAP server daemon (slapd) processed certain search queries requesting only attributes and no values. In certain configurations, a remote attacker could issue a specially-crafted LDAP search query that, when processed by slapd, would cause slapd to crash due to an assertion failure.

Bug Fixes

BZ#784211

When OpenLDAP was set with master-master replication and with the "unique" overlay configured on the back-end database, a server failed to synchronize after getting online. An upstream patch has been applied and the overlay no longer causes breaches in synchronization.

BZ#790687

When the OpenLDAP server was enabled on the ldaps port (636), this port could already be taken by another process using the bindresvport() call. Consequently, the slapd daemon could not bind to the ldaps port. This update adds a configuration file for the portreserve service to reserve the ldaps port and this port is now always available for slapd.

BZ#742163

When the OpenLDAP server was running with the "constraint" overlay enabled and the "count" restrictions configured, specific modify operations could cause "count" restriction violation without the overlay detecting it. Now, the count overlay has been fixed to detect such situations and the server returns the "constraint violation" error as expected.

BZ#783445

If the slapd daemon was set up with master-master replication over TLS, when started, it terminated unexpectedly with a segmentation fault due to accessing unallocated memory. This update applies a patch that copies and stores the TLS initialization parameters, until the deferred TLS initialization takes place and the crashes no longer occur in the described scenario.

BZ#796808

When an OpenLDAP server used TLS and a problem with loading the server key occurred, the server terminated unexpectedly with a segmentation fault due to accessing uninitialized memory. With this update, variables holding TLS certificate and keys are properly initialized, the server no longer crashes in the described scenario, and information about the failure is logged instead.

BZ#807363

Due to a bug in the libldap library, when a remote LDAP server responded with a referral to a client query and the referral chasing was enabled in the library on the client, a memory leak occurred in libldap. An upstream patch has been provided to address this issue and memory leaks no longer occur in the described scenario.

BZ#742023

If a client established a TLS connection to a remote server, which had a certificate issued by a commonly trusted certificate authority (CA), the server certificate was rejected because the CA certificate could not be found. Now, during the package installation, certificate database is created and a module with a trusted root CA is loaded. Trusted CAs shipped with the Mozilla NSS package are used and TLS connections to a remote server now work as expected.

BZ#784203

Under certain conditions, when the unbind operation was called and the ldap handle was destroyed, the library attempted to close the connection socket, which was already closed. Consequently, warning messages from the valgrind utility were returned. An upstream patch has been applied, additional checks before closing a connection socket have been added, and the socket in the described scenario is now closed only once with no warnings returned.

BZ#732916

Previously, description of the SASL_NOCANON option was missing under the "SASL OPTIONS" section in the ldap.conf man page. This update amends the man page.

BZ#743781

When mutually exclusive options "-w" and "-W" were passed to any OpenLDAP client tool, the tool terminated with an assertion error. Upstream patch has been applied and client tools now do not start if these options are passed on the command line together, thus preventing this bug.

BZ#745470

Previously, description of the "-o" and "-N" options was missing in man pages for OpenLDAP client tools. This update amends the man pages.

BZ#730745

When the "memberof" overlay was set up on top of the front end database, the server terminated unexpectedly with a segmentation fault if an entry was modified of deleted. With this update, the "memberof" overlay can no longer be set up on top of the front end database. Instead, it is required to be set up on top the back end database or databases. Now, the crash no longer occurs in the described scenario.

BZ#816168

When a utility from the openldap-clients package was called without a specified URL, a memory leak occurred. An upstream patch has been applied to address this issue and the bug no longer occurs in the described scenario.

BZ#818844

When connecting to a remote LDAP server with TLS enabled, while the TLS_CACERTDIR parameter was set to Mozilla NSS certificate database and the TLS_CACERT parameter was set to PEM bundle with CA certificates, certificates from the PEM bundle were not loaded. If the signing CA certificate was present only in the PEM CA bundle specified by TLS_CACERT, validation of the remote certificate failed. This update allows loading of CA certificates from the PEM bundle file if the Mozilla NSS certificate database is set up as well. As a result, the validation succeeds in the described scenario.

Users of openldap are advised to upgrade to these updated packages, which fix this issue and these bugs.

5.220. OPENMOTIF

5.220.1. RHBA-2012:1405 – openmotif bug fix update

Updated openmotif packages that fix three bugs is now available for Red Hat Enterprise Linux 6.

The openmotif packages include the Motif shared libraries needed to run applications, which are dynamically linked against Motif, as well as MWM, the Motif Window Manager.

Bug Fixes

BZ#866499

Under certain circumstances, closing an application using a text or a combo box widget could result in an invalid memory access for widgets that were already destroyed. A patch has been provided to address this issue and freed memory is no longer accessed in the described scenario.

BZ#866496

Prior to this update, the insertion cursor of a text widget could have a shadow border line under certain settings. This bug has been fixed and the cursor now displays correctly at all times.

BZ#867463

Due to 32-bit time stamp issues, attempting to copy and paste on a 64-bit architecture using the clipboard could fail occasionally. With this update, the underlying source code has been modified to ensure the time stamp always contains a "CARD32" value, so that copy and paste on 64-bit architectures works as expected.

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

5.221. OPENSSH

5.221.1. RHBA-2012:1443 – openssh bug fix update

Updated openssh packages that fix a bug are now available for Red Hat Enterprise Linux 6.

[Updated 12 Nov 2012] This advisory has been updated with an accurate description for BZ#871127 to indicate that the nature of the bug is not architecture-specific. This update does not change the packages in any way.

OpenSSH is OpenBSD's SSH (Secure Shell) protocol implementation. These packages include the core files necessary for both the OpenSSH client and server.

Bug Fix

BZ#871127

When SELinux was disabled on the system, no on-disk policy was installed, an user account was used for a connection, and no "~/.ssh" configuration was present in that user's home directory, the ssh client could terminate unexpectedly with a segmentation fault when attempting to connect to another system. A patch has been provided to address this issue and the crashes no longer occur in the described scenario.

All openssh users are advised to upgrade to these updated packages, which fix this bug.

5.221.2. RHSA-2012:0884 – 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is 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-2011-5000

A denial of service flaw was found in the OpenSSH GSSAPI authentication implementation. A remote, authenticated user could use this flaw to make the OpenSSH server daemon (sshd) use an excessive amount of memory, leading to a denial of service. GSSAPI authentication is enabled by default ("GSSAPIAuthentication yes" in "/etc/ssh/sshd_config").

Bug Fixes

BZ#732955

SSH X11 forwarding failed if IPv6 was enabled and the parameter X11UseLocalhost was set to "no". Consequently, users could not set X forwarding. This update fixes sshd and ssh to correctly bind the port for the IPv6 protocol. As a result, X11 forwarding now works as expected with IPv6.

BZ#744236

The sshd daemon was killed by the OOM killer when running a stress test. Consequently, a user could not log in. With this update, the sshd daemon sets its oom_adj value to -17. As a result, sshd is not chosen by OOM killer and users are able to log in to solve problems with memory.

BZ#809619

If the SSH server is configured with a banner that contains a backslash character, then the client will escape it with another "\" character, so it prints double backslashes. An upstream patch has been applied to correct the problem and the SSH banner is now correctly displayed.

Enhancements

BZ#657378

Previously, SSH allowed multiple ways of authentication of which only one was required for a successful login. SSH can now be set up to require multiple ways of authentication. For example, logging in to an SSH-enabled machine requires both a passphrase and a public key to be entered. The RequiredAuthentications1 and RequiredAuthentications2 options can be configured in the /etc/ssh/sshd_config file to specify authentications that are required for a successful login. For example, to set key and password authentication for SSH version 2, type:

```
echo "RequiredAuthentications2 publickey,password" >> /etc/ssh/sshd_config
```

For more information on the aforementioned /etc/ssh/sshd_config options, refer to the sshd_config man page.

BZ#756929

Previously, OpenSSH could use the Advanced Encryption Standard New Instructions (AES-NI) instruction set only with the AES Cipher-block chaining (CBC) cipher. This update adds support for Counter (CTR) mode encryption in OpenSSH so the AES-NI instruction set can now be used efficiently also with the AES CTR cipher.

BZ#798241

Prior to this update, an unprivileged slave sshd process was run as the sshd_t context during privilege separation (privsep). sshd_t is the SELinux context used for running the sshd daemon. Given that the unprivileged slave process is run under the user's UID, it is fitting to run this process under the user's SELinux context instead of the privileged sshd_t context. With this update, the unprivileged slave process is now run as the user's context instead of the sshd_t context in accordance with the principle of privilege separation. The unprivileged process, which might be potentially more sensitive to security threats, is now run under the user's SELinux context.

Users are advised to upgrade to these updated openssh packages, which contain backported patches to resolve these issues and add these enhancements. After installing this update, the OpenSSH server daemon (sshd) will be restarted automatically.

5.222. OPENSSL

5.222.1. RHBA-2012:1195 – openssl bug fix update

Updated openssl packages that fix a bug are now available for Red Hat Enterprise Linux 6.

OpenSSL is 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 Fix

BZ#848406

When a save operation of a private key file in the encrypted format was attempted in FIPS mode, the resulting file was corrupted because the PEM encryption uses hash algorithms that are not available in FIPS mode. With this update, the PKCS#8 encrypted format is used to write private keys to files in FIPS mode. This file format does not use algorithms unavailable in FIPS mode, thus preventing this bug.

All users of OpenSSL should upgrade to these updated packages, which fix this bug.

5.223. OPENSWAN

5.223.1. RHBA-2012:1305 – openswan bug fix update

Updated openswan packages that fix a bug are now available for Red Hat Enterprise Linux 6.

Openswan is a free implementation of IPsec (Internet Protocol Security) and IKE (Internet Key Exchange) for Linux. The openswan packages contain daemons and user-space tools for setting up Openswan. It supports the NETKEY/XFRM IPsec kernel stack that exists in the default Linux kernel. Openswan 2.6 and later also supports IKEv2 (Internet Key Exchange Protocol version 2), which is defined in RFC5996.

Bug Fix

BZ#852454

When a tunnel was established between two IPsec hosts (say host1 and host2) utilizing DPD (Dead Peer Detection), 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 SA (Security Association), 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.

All users of openswan are advised to upgrade to these updated packages, which fix this bug.

5.223.2. RHBA-2012:1069 – openswan bug fix update

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

Openswan is a free implementation of IPsec (internet Protocol Security) and IKE (Internet Key Exchange) for Linux. The openswan packages contain the daemons and user-space tools for setting up Openswan. It supports the NETKEY/XFRM IPsec kernel stack that exists in the default Linux kernel. Openswan 2.6 and later also supports IKEv2 (Internet Key Exchange Protocol Version 2), which is defined in RFC5996.

Bug Fixes

BZ#834660

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 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#834662

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 setup in transport mode.

All users of openswan are advised to upgrade to these updated packages, which fix these bugs.

5.223.3. RHBA-2012:0916 – openswan bug fix update

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

Openswan is a free implementation of IPsec (Internet Protocol Security) and IKE (Internet Key Exchange) for Linux. The openswan package contains the daemons and user space tools for setting up Openswan. It supports the NETKEY/XFRM IPsec kernel stack that exists in the default Linux kernel. Openswan 2.6.x also supports IKEv2 (*RFC4306*).

Bug Fixes

BZ#768162

Previously, **Openswan** sometimes generated a KE payload that was 1 byte shorter than specified by the Diffie-Hellman algorithm. Consequently, IKE renegotiation failed at random intervals. An error message in the following format was logged:

> next payload type of ISAKMP Identification Payload has an unknown value:

This update checks the length of the generated key and if it is shorter than required, leading zero bytes are added.

BZ#768442

Older versions of kernel required the output length of the HMAC hash function to be truncated to 96 bits, therefore **Openswan** previously worked with 96-bit truncation length when using the HMAC-SHA2-256 algorithm. However, newer kernels require the 128-bit HMAC truncation length, which is as per the *RFC4868* specification. Consequently, this difference could cause incompatible SAs to be set on IKE endpoints due to one endpoint using 96-bit and the other 128-bit output length of the hash function. This update modifies the underlying code so that Openswan now complies with *RFC4868* and adds support for the new kernel configuration parameter, sha2_truncbug. If the ***sha2_truncbug*** parameter is set to **yes**, Openswan now passes the correct key length to the kernel, which ensures interoperability between older and newer kernels.

BZ#771457

When processing an IKE_SA_INIT exchange and the RESERVED field of the IKE_SA_INIT request or response messages was modified, **Openswan** did not ignore the field as expected according to the IKEv2 *RFC5996* specification. Consequently, IKE_SA_INIT messages with reserved fields set were processed as erroneous messages by Openswan and the IKE_SA_INIT exchange failed. With this update, Openswan has been modified to ignore reserved fields as expected and IKE_SA_INIT exchanges succeed in this scenario.

BZ#771460

When processing an IKE_AUTH exchange and the RESERVED field of the IKE_AUTH request or response messages was modified, **Openswan** did not ignore the field as expected according to the IKEv2 *RFC5996* specification. Consequently, the IKE_AUTH messages were processed as erroneous messages by Openswan and the IKE_AUTH exchange failed. With this update, Openswan has been modified to ignore reserved fields as expected and IKE_AUTH exchanges succeed in this scenario.

BZ#771461

Openswan incorrectly processed traffic selector messages proposed by the responder (the endpoint responding to an initiated exchange) by failing to confine them to a subset of the initially

proposed traffic selectors. As a consequence, Openswan set up CHILD security associations (SAs) incorrectly. With this update, Openswan reduces the set of traffic selectors correctly, and sets up IKE CHILD SAs accordingly.

BZ#771463

Previously, **Openswan** did not behave in accordance with the IKEv2 *RFC5996* specification and ignored IKE_AUTH messages that contained an unrecognized "Notify" payload. This resulted in IKE SAs being set up successfully. With this update, Openswan processes any unrecognized Notify payload as an error and IKE SA setup fails as expected.

BZ#771464

When processing an INFORMATIONAL exchange, the responder previously did not send an INFORMATIONAL response message as expected in reaction to the INFORMATIONAL request message sent by the initiator. As a consequence, the INFORMATIONAL exchange failed. This update corrects **Openswan** so that the responder now sends an INFORMATIONAL response message after every INFORMATIONAL request message received, and the INFORMATIONAL exchange succeeds as expected in this scenario.

BZ#771465

When processing an INFORMATIONAL exchange with a Delete payload, the responder previously did not send an INFORMATIONAL response message as expected in reaction to the INFORMATIONAL request message sent by the initiator. As a consequence, the INFORMATIONAL exchange failed and the initiator did not delete IKE SAs. This updates corrects **Openswan** so that the responder now sends an INFORMATIONAL response message and the initiator deletes IKE SAs as expected in this scenario.

BZ#771466

When the responder received an INFORMATIONAL request with a "Delete" payload for a CHILD SA, **Openswan** did not process the request correctly and did not send the INFORMATIONAL response message to the initiator as expected according to the *RFC5996* specification. Consequently, the responder was not aware of the request and only the initiator's CHILD SA was deleted. With this update, Openswan sends the response message as expected and the CHILD SA is deleted properly on both endpoints.

BZ#771467

Openswan did not ignore the minor version number of the IKE_SA_INIT request messages as required by the *RFC5996* specification. Consequently, if the minor version number of the request was higher than the minor version number of the IKE protocol used by the receiving peer, Openswan processed the IKE_SA_INIT messages as erroneous and the IKE_SA_INIT exchange failed. With this update, Openswan has been modified to ignore the Minor Version fields of the IKE_SA_INIT requests as expected and the IKE_SA_INIT exchange succeeds in this scenario.

BZ#771470

The **Openswan** IKEv2 implementation did not correctly process an IKE_SA_INIT message containing an INVALID_KE_PAYLOAD "Notify" payload. With this fix, Openswan now sends the INVALID_KE_PAYLOAD notify message back to the peer so that IKE_SA_INIT can restart with the correct KE payload.

BZ#771472

Openswan incorrectly processed traffic selector messages proposed by the initiator (the endpoint which started an exchange) by failing to confine them to a subset of the initially proposed traffic selectors. As a consequence, Openswan set up CHILD SAs incorrectly. With this update, Openswan

reduces the set of traffic selectors correctly, and sets up IKE CHILD SAs accordingly.

BZ#771473

Previously, **Openswan** did not respond to INFORMATIONAL requests with no payloads that are used for dead-peer detection. Consequently, the initiator considered the responder to be a dead peer and deleted the respective IKE SAs. This update modifies Openswan so that an empty INFORMATIONAL response message is now sent to the initiator as expected, and the initiator no longer incorrectly deletes IKE SAs in this scenario.

BZ#771475

When processing an INFORMATIONAL exchange and the RESERVED field of the INFORMATIONAL request or response messages was modified, **Openswan** did not ignore the field as expected according to the IKEv2 *RFC5996* specification. Consequently, the INFORMATIONAL messages were processed as erroneous by Openswan, and the INFORMATIONAL exchange failed. With this update, Openswan has been modified to ignore reserved fields as expected and INFORMATIONAL exchanges succeed in this scenario.

BZ#795842

When the initiator received an INFORMATIONAL request with a "Delete" payload for an IKE SA, **Openswan** did not process the request correctly and did not send the INFORMATIONAL response message to the responder as expected according to the *RFC5996* specification. Consequently, the initiator was not aware of the request and only the responder's IKE SA was deleted. With this update, Openswan sends the response message as expected and the IKE SA is deleted properly on both endpoints.

BZ#795850

IKEv2 requires each IKE message to have a sequence number for matching a request and response when re-transmitting the message during the IKE exchange. Previously, **Openswan** incremented sequence numbers incorrectly so that IKE messages were processed in the wrong order. As a consequence, any messages sent by the responder were not processed correctly and any subsequent exchange failed. This update modifies Openswan to increment sequence numbers in accordance with the *RFC5996* specification so that IKE messages are matched correctly and exchanges succeed as expected in this scenario.

Users of openswan should upgrade to this updated package, which fixes these bugs.

5.223.4. RHBA-2013:1161 – openswan bug fix update

Updated openswan packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

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 Fix

BZ#983451

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.

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

5.224. OPROFILE

5.224.1. RHBA-2012:0966 – oprofile bug fix and enhancement update

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

OProfile is a system-wide profiler for Linux systems. The profiling runs transparently in the background and profile data can be collected at any time. OProfile uses the hardware performance counters provided on many processors, and can use the Real Time Clock (RTC) for profiling on processors without counters.

The oprofile packages have been upgraded to upstream version 0.9.7, which provides a number of bug fixes and enhancements over the previous version. (BZ#739142)

Bug Fix

BZ#748789

Under certain circumstances, the "opannotate" and "opreport" commands reported no results. With this update, this problem has been fixed so that these commands work as expected.

All OProfile users are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.225. ORBIT2

5.225.1. RHBA-2012:1457 – ORBit2 bug fix update

Updated ORBit2 packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The ORBit2 packages provide a high-performance Object Request Broker (ORB) for Common Object Request Broker Architecture (CORBA). ORBit allows programs to send requests and receive replies from other programs, regardless of where the programs are located. CORBA is a standard that enables communication between program objects, regardless of the programming language and platform used.

Bug Fix

BZ#866469

ORBit2 connection code was unable to process the EAGAIN error, which can be returned by an AF_UNIX socket due to too many simultaneous connection requests. This could cause gconfd-2 clients to fail to connect to the gconfd daemon. Consequently, various errors were emitted when logging in to the GNOME 2 session and the server was under a heavy load. This update adds the patch that enables the ORBit2 request broker to wait for the defined time period if the EAGAIN error is received, and then retry to establish the connection. Errors no longer occur when logging in to the GNOME 2 session in this scenario.

All users of ORBit2 are advised to upgrade to these updated packages, which fix this bug.

5.226. PACEMAKER

5.226.1. RHBA-2012:0846 – pacemaker bug fix and enhancement update

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

Pacemaker is a high-availability cluster resource manager with a powerful policy engine.

Bug Fixes

BZ#720214

Previously, the "port" parameter was declared as required in the fencing agent metadata. Some devices, however, determine the correct port automatically and therefore do not need it specified in their configuration. As a consequence, updating configurations that didn't contain the "port" parameter failed with the following error message:

```
ERROR: apc-fencing: required parameter port not defined
```

This update modifies the fencing agent metadata so that "port" is not declared as a required configuration option. As a result, updating configurations in which the "port" parameter is not specified succeeds.

BZ#720218

Previously, the "start" and "stop" actions were not declared in the fencing agent metadata. As a consequence, the following warning messages were displayed when configuring fencing devices:

```
WARNING: apc-fencing: action start not advertised in meta-data, it may not be supported by the RA
WARNING: apc-fencing: action stop not advertised in meta-data, it may not be supported by the RA
```

This update adds the "start" and "stop" actions to all fencing agent metadata. As a result, configuring fencing devices no longer displays the aforementioned warning messages.

BZ#789397

Due to an error in the underlying source code, operation failure records were not removed together with removed resources. Consequently, resources re-defined after previous deletion were associated with the previous failure records. This update modifies the underlying source code so that failure records are removed together with removed resources. As a consequence, previously deleted resources are not associated with any failure records after being re-defined.

BZ#799070

Previously, the logic for determining whether a resource is active was incorrect. Consequently, active resources on nodes in the "UNCLEAN" state were ignored by tools that relied on this logic. This update fixes the logic. As a result, tools relying on this logic report active resources on nodes in the "UNCLEAN" state as active.

BZ#801351

Previously, descriptions of the -v and -V options of the crm_report command were swapped in text of its manual pages. Consequently, some users were mislead by unexpected behavior when using these options. This update corrects the text of the manual pages so that they reflect the actual behavior of the crm_report command.

Enhancement

BZ#782255

Pacemaker now uses the libqb library for logging. This provides less verbose logs while still providing the ability to debug and support Pacemaker.

All users of Pacemaker are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.227. PACKAGEKIT

5.227.1. RHBA-2012:0913 – PackageKit bug fix update

Updated PackageKit packages that fix four bugs are now available for Red Hat Enterprise Linux 6.

PackageKit is a D-Bus abstraction layer that allows the session user to manage packages in a secure way using a cross-distribution, cross-architecture API.

Bug Fixes

BZ#744359

If a user attempted to install or update packages that were either unsigned, or signed with a GPG key that was not installed and not available for installation, PackageKit kept reporting that the packages are from an untrusted source and repeatedly prompted the user for confirmation. This update corrects this behavior and ensures that up-to-date interfaces are used to specify that untrusted packages are to be installed.

BZ#783537

If the pkcon console client is executed with the "--noninteractive" command line option, it is not supposed to prompt the user for any confirmation. Previously, running pkcon with this option did not prevent it from requiring confirmation in certain situations, such as if a package signing key had to be imported or additional dependent packages needed to be removed during package removal. With this update, the pkcon utility no longer requires confirmation in these situations and proceeds as if the user answered "yes" on the command line.

BZ#684861, BZ#700448

Prior to this update, the PackageKit-yum and PackageKit-yum-plugin subpackages did not specify pygobject2 and dbus-python as their dependencies. Consequently, if these packages were not present on the system, certain tools such as the pkcon console client or the refresh-packagekit plug-in for YUM did not work properly. This update adapts the PackageKit-yum and PackageKit-yum-plugin subpackages to require pygobject2 and dbus-python respectively. As a result, the tools that are installed with either of these subpackages no longer fail to work properly due to missing dependencies.

All users of PackageKit are advised to upgrade to these updated packages, which fix these bugs.

5.228. PAM_PKCS11

5.228.1. RHBA-2012:0972 – pam_pkcs11 bug fix update

Updated pam_pkcs11 packages which fix a bug are now available for Red Hat Enterprise Linux 6.

The pam_pkcs11 package allows X.509 certificate-based user authentication. It provides access to the certificate and its dedicated private key with an appropriate Public Key Cryptographic Standards #11 (PKCS#11) module.

Bug Fix

BZ#756917

When remotely logged into a system with smart card log-in turned on, users saw the following unnecessary error message when trying to use su:

```
ERROR:pam_pkcs11.c:224: Remote login (from localhost:13.0) is not (yet) supported
```

The user was still able to use su despite this message. With this update the message is logged but no longer displayed.

All users of pam_pkcs11 are advised to upgrade to these updated packages, which fix this bug.

5.229. PANGO

5.229.1. RHBA-2012:1498 – pango bug fix update

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

Pango is a library for laying out and rendering of text, with an emphasis on internationalization. Pango forms the core of text and font handling for GTK+ widget toolkit.

Bug Fix

BZ#878772

Due to a regression in the pangoft2 module, an incorrect calculation caused that the approximate width of a character was reported as zero. Consequently, some applications using the Pango library failed with an assertion. With this update, the approximate character width is always set correctly and the failures no longer occur.

User of pango should upgrade to these updated packages, which fix this bug.

5.230. PARTED

5.230.1. RHBA-2012:0773 – parted bug fix update

Updated parted packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The parted packages allow you 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#698121, BZ#751164

Prior to this update, editing partitions on a mpath device udev could, under certain circumstances, interfere with re-reading the partition table. This update adds the dm_udev_wait option so that udev now correctly synchronizes.

BZ#750395

Prior to this update, the libparted partition_duplicate() function did not correctly copy all GPT partition flags. This update modifies the underlying code so that all flags are correctly copied and adds a test to ensure correct operation.

All parted users are advised to upgrade to these updated packages, which fix these bugs.

5.231. PCRE

5.231.1. RHBA-2012:1240 – pcre bug fix release

Updated pcre packages that fix four bugs are now available for Red Hat Enterprise Linux 6.

The pcre packages provide the Perl-compatible regular expression (PCRE) library.

Bug Fixes

BZ#756105

Prior to this update, matching patterns with repeated forward reference failed to match if the first character was not repeated at the start of the matching text. This update modifies the matching algorithm not to expect the first character again. Now, patterns with repeated forward references match as expected.

BZ#759475

Prior to this update, case-less patterns in UTF-8 mode did not match characters at the end of input text with encoding length that was shorter than the encoding length of character in the pattern, for example "/ȁ/8i".This update modifies the pcre library to count the length of matched characters correctly. Now, case-less patterns match characters with different encoding length correctly even at the end of an input string.

BZ#799003

Prior to this update, manual pages for the pcre library contained misprints. This update modifies the manual pages.

BZ#842000

Prior to this update, applications that were compiled with the libpcrecpp library from the pcre version 6 could not been executed against libpcrecpp library from the pcre version 7 because the application binary interface (ABI) was mismatched. This update adds the compat RE::Init() function for the pcre version 6 to the pcre version 7 libpcrecpp library. Applications that were compiled on Red Hat Enterprise Linux 5 and use the RE::Init function can now be executed on Red Hat Enterprise Linux 6.

All users of pcre are advised to upgrade to these updated packages, which fix these bugs.

5.231.2. RHBA-2012:0445 – pcre bug fix update

Updated pcre packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The pcre packages provide the Perl-compatible regular expression (PCRE) library.

Bug Fixes

BZ#613690

Prior to this update, the PCRE license file was missing from the pcre-static subpackage. This update adds the license file.

BZ#676636

Prior to this update, the pcretest(1) manual page contained various misprints. This update corrects these misprints.

BZ#676643

Prior to this update, the pcregrep(1) manual page contained various misprints. This update corrects these misprints.

All users of pcre are advised to upgrade to these updated packages, which fix these bugs.

5.232. PCSC-LITE

5.232.1. RHBA-2012:1343 – pcsc-lite bug fix update

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

PC/SC Lite provides a Windows SCard compatible interface for communicating with smart cards, smart card readers, and other security tokens.

Bug Fix

BZ#851199

Despite the update described in the RHBA-2012:0990 advisory, the chkconfig utility did not automatically place the pcscd init script after the start of the HAL daemon. Consequently, pcscd was unable to recognize USB readers. With this update, the pcscd init script has been changed to explicitly start only after HAL is up, thus fixing this bug.

All pcsc-lite users are advised to upgrade to these updated packages, which fix this bug.

5.232.2. RHBA-2012:0990 – pcsc-lite bug fix update

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

PC/SC Lite provides a Windows SCard compatible interface for communicating with smart cards, smart card readers, and other security tokens.

Bug Fix

BZ#812469

Previously, the pcscd init script pointed to the wrong value to identify the HAL daemon. It also wrongly started at runlevel 2. As a result, chkconfig did not automatically place pcscd after the start of the HAL daemon, thus pcscd failed to see USB readers. With this update, the pcscd init script has been changed to properly identify the HAL daemon and only start at runlevels 3, 4, and 5, thus fixing this bug.

All pcsc-lite users are advised to upgrade to these updated packages, which fix this bug.

5.233. PERL-DBD-PG

5.233.1. RHSA-2012:1116 – Moderate: perl-DBD-Pg security update

An updated perl-DBD-Pg package that fixes two security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Perl DBI is a database access Application Programming Interface (API) for the Perl language. perl-DBD-Pg allows Perl applications to access PostgreSQL database servers.

Security Fix

CVE-2012-1151

Two format string flaws were found in perl-DBD-Pg. A specially-crafted database warning or error message from a server could cause an application using perl-DBD-Pg to crash or, potentially, execute arbitrary code with the privileges of the user running the application.

All users of perl-DBD-Pg are advised to upgrade to this updated package, which contains a backported patch to fix these issues. Applications using perl-DBD-Pg must be restarted for the update to take effect.

5.234. PERL-GSSAPI

5.234.1. RHBA-2012:1340 – perl-GSSAPI bug fix update

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

The perl-GSSAPI packages provide Perl extension for GSSAPIv2 access.

Bug Fix

BZ#657274

Prior to this update, the perl-GSSAPI specification file used a krb5-devel file which was removed. As a consequence, the perl-GSSAPI package could not be rebuilt. This update modifies the specification file to use the current krb5-devel files.

All users of perl-GSSAPI are advised to upgrade to these updated packages, which fix this bug.

5.235. PERL-IPC-RUN3

5.235.1. RHBA-2012:1440 – perl-IPC-Run3 bug fix update

Updated perl-IPC-Run3 packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The perl-IPC-Run3 packages provide a module to run subprocesses and redirect the stdin, stdout, and stderr functionalities to files and perl data structures. The perl-IPC-Run3 package allows to use system, qx, and open3 modules with a simple API.

Bug Fixes

BZ#657487

Prior to this update, binary perl-IPC-Run3 packages failed to build if the perl-Time-HiRes module was not installed. This update adds the perl-Time-HiRes package to the build-time dependencies for perl-IPC-Run3.

BZ#870089

Prior to this update, tests that called the IP-Run3 profiler failed when the internal perl-IPC-Run3 test suite was used. This update, adds run-time dependencies on perl(Getopt::Long) and perl(Time::HiRes) to the perl-IPC-Run3 package because certain IP-Run3 functions require the perl modules. Now, the IPC-Run3 profiler runs as expected.

All users of perl-IPC-Run3 are advised to upgrade to these updated packages, which fix these bugs.

5.236. PERL-IPC-RUN

5.236.1. RHBA-2012:1336 – perl-IPC-Run bug fix update

Updated perl-IPC-Run packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The perl-IPC-Run packages provide a mechanism for Perl scripts to interact with child processes.

Bug Fix

BZ#856840

Prior to this update, the IO::Pty Perl module was not loaded when using the command "IPC::Run::harness" with the ">pty>" argument if the perl-IO-Tty package was not installed. As a consequence, the Perl code failed. This update adds a perl-IO-Tty dependency to the perl-IPC-Run packages.

All users of perl-IPC-Run are advised to upgrade to these updated packages, which fix this bug.

5.237. PERL-SOAP-LITE

5.237.1. RHBA-2012:1388 – perl-SOAP-Lite bug fix update

An updated perl-SOAP-Lite package that fixes one bug is now available for Red Hat Enterprise Linux 6.

SOAP::Lite is a collection of Perl modules, which provides a simple and lightweight interface to the Simple Object Access Protocol (SOAP) both on client and server side.

Bug Fix

BZ#748376

XMLRPC requests could fail if the MOD_PERL environment value was defined. The standard read() function is now used instead of the sysread() function when MOD_PERL is defined. As a result, XMLRPC no longer fails in this scenario.

All users of perl-SOAP-Lite are advised to upgrade to this updated package, which fixes this bug.

5.238. PERL-SYS-VIRT

5.238.1. RHBA-2012:0754 – perl-Sys-Virt bug fix and enhancement update

Updated perl-Sys-Virt packages that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The perl-Sys-Virt packages provide application programming interfaces (APIs) to manage virtual machines from Perl with the libvirt library.

The perl-Sys-Virt package has been upgraded to upstream version 0.9.10, which provides a number of bug fixes and enhancements over the previous version. (BZ#752436)

Bug Fixes

BZ#661801

Prior to this update, the perl-Sys-Virt spec file did not contain the "perl(Time::HiRes)" requirement. As a consequence, perl-Sys-Virt could not be rebuild in mock mode. This update adds the missing requirement to the spec file. Now, perl-Sys-Virt can is rebuild in mock mode as expected.

BZ#747483

Prior to this update, the perl-Sys-Virt man page did not document the "$flags" parameter for the "get_xml_description" executable. This update modifies the man page so that the parameter is correctly documented.

BZ#748689

Prior to this update, the default settings for the remote domain memory statistics used a length of 16 bits only. As a consequence, the get_node_cpu_stats() function could send the libvirt error "code: 1, message: internal error nparams too large". This update modifies libvirt so that the maximum length is now 1024 bits.

BZ#773572

Prior to this update, the bandwidth in the "block_pull and set_block_job_speed" method was incorrectly given in Kilobytes per second (Kb/s). This update changes the bandwidth unit to Megabytes per second (Mb/s).

BZ#800766

Prior to this update, the bandwidth for maximum migration bandwidth was incorrectly given in Kilobytes per second (Kb/s). This update changes the maximum migration bandwidth unit to Megabytes per second (Mb/s).

BZ#809906

Prior to this update, the documentation for "Sys::Virt::StoragePool" incorrectly stated that the object method "get_info()" returns a hash. This update corrects this misprint and correctly states that object method returns a hash reference.

Enhancement

BZ#800734

Prior to this update, the Perl API bindings could not handle tunable parameters in string format. As a consequence, the block I/O tunable parameters could not be read or updated. This update adds support for string parameters. Now, the block I/O tunable parameters can be read and updated from the Perl API.

All users of perl-Sys-Virt are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.239. PERL

5.239.1. RHBA-2012:0843 – perl bug fix and enhancement update

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

Perl is a high-level programming language commonly used for system administration utilities and web programming.

Bug Fixes

BZ#707960, BZ#717565

Previously, the perl and perl-libs packages owned the /usr/local/lib/perl5/, /usr/local/lib64/perl5/, and /usr/local/share/perl5/ directories. Consequent to this, when the /usr/local/ directory was read-only or contained local changes, the perl and perl-libs packages could not be installed, reinstalled, updated, or removed. With this update, the perl and perl-libs packages have been updated not to own the aforementioned directories. This ensures that after applying this update, both these packages can be installed, reinstalled, updated, or removed as expected and that such an action no longer affects the /usr/local/ directory.

BZ#738421

Prior to this update, an attempt to use bzip2 compression or decompression in a Perl program (for example, by using the IO::Compress::Bzip2 module) failed with an error, because no bzip2 module was available. This update adds support for bzip2 compression and decompression, and provides two new packages, perl-Compress-Raw-Bzip2 and perl-IO-Compress-Bzip2, for the IO:Compress::Bzip2 and Compress::Raw::Bzip2 modules respectively.

BZ#750145

When the perl-ExtUtils-MakeMaker package was not present on the system, an attempt to run the cpanp utility that is provided by the perl-CPANPLUS package previously failed with the following error:

```
Can't locate ExtUtils/MakeMaker.pm in @INC
```

With this update, perl-ExtUtils-MakeMaker is now required by the perl-IPC-Cmd package, and the cpanp utility installed from the RPM package now works as expected.

BZ#801804

Due to an error in the corresponding spec file, the version of the perl-Compress-Raw-Zlib package was higher than the version reported by the Compress::Raw::Zlib module it provides. This update changes the package version to "2.020" so that it matches the version of the Perl module. Additionally, this update changes the epoch number to preserve the RPM package version string ordering.

BZ#805606

Prior to this update, when a POSIX::strftime() function call returned a string that was longer than 64 bytes, a memory leak occurred. Consequent to this, any script that repeatedly called this function could consume a significant amount of memory over time. This update applies an upstream patch that adapts the implementation of the POSIX::strftime() function to reallocate the memory instead of allocating new one, and such memory leaks no longer occur.

BZ#806373

When the nextStream() method was called on an IO::Uncompress::Unzip object with the last stream in the ZIP decoder, it did not behave according to the documentation and returned a non-zero value. With this update, an upstream patch has been applied to correct this error, and the nextStream() method now behaves as documented.

Enhancement

BZ#817480

Previous versions of the Perl interpreter were not compiled with the "usesitecustomize" feature, which rendered users unable to use the /usr/local/share/perl5/sitecustomize.pl script to automatically modify the Perl environment when the interpreter is executed. With this update, the Perl interpreter has been recompiled with the "-Dusesitecustomize" option so that the interpreter now automatically executes the /usr/local/share/perl5/sitecustomize.pl script before interpreting the Perl code.

All users of perl are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.240. PHP-PECL-APC

5.240.1. RHSA-2012:0811 – Low: php-pecl-apc security, bug fix, and enhancement update

Updated php-pecl-apc packages that fix one security issue, several bugs, and add various enhancements are now available for Red Hat Enterprise Linux 6.

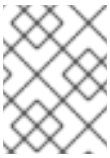
The Red Hat Security Response Team has rated this update as having low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

The php-pecl-apc packages contain APC (Alternative PHP Cache), the framework for caching and optimization of intermediate PHP code.

Security Fix

CVE-2010-3294

A cross-site scripting (XSS) flaw was found in the "apc.php" script, which provides a detailed analysis of the internal workings of APC and is shipped as part of the APC extension documentation. A remote attacker could possibly use this flaw to conduct a cross-site scripting attack.



NOTE

The administrative script is not deployed upon package installation. It must manually be copied to the web root (the default is "/var/www/html/", for example).

In addition, the php-pecl-apc packages have been upgraded to upstream version 3.1.9, which provides a number of bug fixes and enhancements over the previous version. (BZ#662655)

All users of php-pecl-apc are advised to upgrade to these updated packages, which fix these issues and add these enhancements. If the "apc.php" script was previously deployed in the web root, it must manually be re-deployed to replace the vulnerable version to resolve this issue.

5.241. PHP-PECL-MEMCACHE

5.241.1. RHBA-2012:0927 – php-pecl-memcache bug fix update

Updated php-pecl-memcache packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The php-pecl-memcache packages enable PHP scripts to use the memcached caching daemon.

Bug Fix

BZ#722418

A bug in memory handling could cause memory corruption when converting a serialized object to a PHP object. Consequently, a PHP script terminated unexpectedly with a segmentation fault. This update corrects the bug so that memory corruption no longer occurs and PHP scripts are now executed successfully.

All users of php-pecl-memcache are advised to upgrade to these updated packages, which fix this bug.

5.242. PHP

5.242.1. RHSA-2012:1046 – Moderate: php security update

Updated php packages that fix multiple security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

PHP is an HTML-embedded scripting language commonly used with the Apache HTTP Server.

Security Fixes

CVE-2012-0057

It was discovered that the PHP XSL extension did not restrict the file writing capability of libxslt. A remote attacker could use this flaw to create or overwrite an arbitrary file that is writable by the user running PHP, if a PHP script processed untrusted eXtensible Style Sheet Language Transformations (XSLT) content.

CVE-2012-1172

Note: This update disables file writing by default. A new PHP configuration directive, "xsl.security_prefs", can be used to enable file writing in XSLT.

A flaw was found in the way PHP validated file names in file upload requests. A remote attacker could possibly use this flaw to bypass the sanitization of the uploaded file names, and cause a PHP script to store the uploaded file in an unexpected directory, by using a directory traversal attack.

CVE-2012-2386

Multiple integer overflow flaws, leading to heap-based buffer overflows, were found in the way the PHP phar extension processed certain fields of tar archive files. A remote attacker could provide a specially-crafted tar archive file that, when processed by a PHP application using the phar extension, could cause the application to crash or, potentially, execute arbitrary code with the privileges of the user running PHP.

CVE-2010-2950

A format string flaw was found in the way the PHP phar extension processed certain PHAR files. A remote attacker could provide a specially-crafted PHAR file, which once processed in a PHP application using the phar extension, could lead to information disclosure and possibly arbitrary code execution via a crafted phar:// URI.

CVE-2012-2143

A flaw was found in the DES algorithm implementation in the crypt() password hashing function in PHP. If the password string to be hashed contained certain characters, the remainder of the string was ignored when calculating the hash, significantly reducing the password strength.

CVE-2012-2336

Note: With this update, passwords are no longer truncated when performing DES hashing. Therefore, new hashes of the affected passwords will not match stored hashes generated using vulnerable PHP versions, and will need to be updated.

It was discovered that the fix for CVE-2012-1823, released via RHSA-2012:0546, did not properly filter all php-cgi command line arguments. A specially-crafted request to a PHP script could cause the PHP interpreter to execute the script in a loop, or output usage information that triggers an Internal Server Error.

CVE-2012-0789

A memory leak flaw was found in the PHP strtotime() function call. A remote attacker could possibly use this flaw to cause excessive memory consumption by triggering many strtotime() function calls.

CVE-2012-0781

A NULL pointer dereference flaw was found in the PHP tidy_diagnose() function. A remote attacker could use specially-crafted input to crash an application that uses tidy::diagnose.

CVE-2011-4153

It was found that PHP did not check the zend_strndup() function's return value in certain cases. A remote attacker could possibly use this flaw to crash a PHP application.

Upstream acknowledges Rubin Xu and Joseph Bonneau as the original reporters of CVE-2012-2143.

All php users should upgrade to these updated packages, which contain backported patches to resolve these issues. After installing the updated packages, the httpd daemon must be restarted for the update to take effect.

5.242.2. RHSA-2013:1061 – Critical: php security update

Updated php packages that fix one security issue are now available for Red Hat Enterprise Linux 5.3 Long Life, and Red Hat Enterprise Linux 5.6, 6.2 and 6.3 Extended Update Support.

The Red Hat Security Response Team has rated this update as having critical 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.

PHP is an HTML-embedded scripting language commonly used with the Apache HTTP Server.

Security Fix

CVE-2013-4113

A buffer overflow flaw was found in the way PHP parsed deeply nested XML documents. If a PHP application used the xml_parse_into_struct() function to parse untrusted XML content, an attacker able to supply specially-crafted XML could use this flaw to crash the application or, possibly, execute arbitrary code with the privileges of the user running the PHP interpreter.

All php users should upgrade to these updated packages, which contain a backported patch to resolve this issue. After installing the updated packages, the httpd daemon must be restarted for the update to take effect.

5.243. PIDGIN

5.243.1. RHSA-2012:1102 – Moderate: pidgin security update

Updated pidgin packages that fix three security issues are now available for Red Hat Enterprise Linux 5 and 6.

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

Pidgin is an instant messaging program which can log in to multiple accounts on multiple instant messaging networks simultaneously.

Security Fixes

CVE-2012-1178

A flaw was found in the way the Pidgin MSN protocol plug-in processed text that was not encoded in UTF-8. A remote attacker could use this flaw to crash Pidgin by sending a specially-crafted MSN message.

CVE-2012-2318

An input validation flaw was found in the way the Pidgin MSN protocol plug-in handled MSN notification messages. A malicious server or a remote attacker could use this flaw to crash Pidgin by sending a specially-crafted MSN notification message.

CVE-2012-3374

A buffer overflow flaw was found in the Pidgin MXit protocol plug-in. A remote attacker could use this flaw to crash Pidgin by sending a MXit message containing specially-crafted emoticon tags.

Red Hat would like to thank the Pidgin project for reporting the CVE-2012-3374 issue. Upstream acknowledges Ulf Härnhammar as the original reporter of CVE-2012-3374.

All Pidgin users should upgrade to these updated packages, which contain backported patches to resolve these issues. Pidgin must be restarted for this update to take effect.

5.244. PIRANHA

5.244.1. RHBA-2012:0891 – piranha bug fix update

An updated piranha package that fixes multiple bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

Piranha provides high-availability and load-balancing services for Red Hat Enterprise Linux. The piranha package contains 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#747300

Previously, the pulse service did not properly close the configuration file. This caused file descriptors leaks in the pulse service and could potentially trigger SELinux AVC errors. With this update, the configuration file is properly closed after reading, and no SELinux errors are observed under these circumstances.

BZ#749594

Previously, the pulse service did not correctly stop the ipvsadm sync daemon due to incorrect ipvsadm syntax. As a consequence, multiple sync daemons existed after restarting the pulse service. With this update, the correct syntax is used. The pulse service now stops all the sync daemons, and exactly one master sync daemon and one backup sync daemon exist at any given time.

BZ#785720

Previously, the lvsd daemon did not correctly identify the existence of a new virtual server when re-reading the configuration file. As a consequence, the lvsd daemon could terminate unexpectedly with a segmentation fault when the pulse service was reloaded. With this update, the lvsd daemon correctly determines if a virtual server has been added to the configuration file when the pulse service is reloaded.

BZ#798362

Previously, the pulse init script did not properly format the output when running the "service pulse reload" command. This update fixes this formatting error by printing a newline character after the init script completes the reload command.

BZ#813906

Previously, the pulse daemon did not correctly detect when the lvsd daemon had been terminated. As a result, the pulse daemon did not trigger a failover. With this update, the pulse daemon correctly detects when lvsd has been terminated and, if a backup director is configured and active, will result in a failover.

BZ#815887

Previously, nanny processes did not correctly write messages to the system log when a sorry_server was defined for a virtual service. When a sorry_server is configured, all nanny processes are run with the "--nodaemon" option. This option prevented messages from being written to syslog and therefore the /var/log/messages file did not contain information about nanny processes. With this update, all nanny processes will write messages to the system log.

Enhancements

BZ#717556

This update adds the ability to specify the sync ID to be used with the ipvsadm sync daemon. A new option in the lvs.cf file, "syncd_id", can be used to set the sync ID. This option can also be configured in the Piranha web interface under the "Redundancy" tab. The default value is 0.

BZ#745271

This update adds the ability to specify IPVS timeouts including TCP session timeout, TCP FIN session timeout, and UDP packet timeout. Three new options have been added to the lvs.cf file: "tcp_timeout", "tcpfin_timeout", and "udp_timeout". These timeout values can also be configured in the Piranha web interface under "Global Settings". The default value for each timeout is 0, which causes no changes to be made to the existing timeouts.

BZ#788541

This update adds the ability to specify the network interface that the ipvsadm sync daemon will use to send and receive multicast messages. A new option in the lvs.cf file, "syncd_iface", can be used to set the sync daemon interface. This option can also be configured in the Piranha web interface under the "Redundancy" tab. The default value is "eth0".

All users of piranha are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.245. PKI-CORE

5.245.1. RHBA-2012:0761 – pki-core bug fix update

Updated pki-core packages that fix several 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 the Red Hat Enterprise Identity (IPA).

Bug Fixes

BZ#745677

A Firefox launcher setting which opened a non-functional Certificate Authority (CA) page was improperly created and applied to all user profiles. With this update, all PKI-related desktop icons have been removed and the problem no longer occurs.

BZ#769388

The pkisilent script did not accept special shell characters, such as spaces or quotation marks, in argument values even if they were properly escaped. Consequently, errors occurred and the script failed. This update improves the code and the problem no longer occurs.

BZ#771790

When installing IPA, the installer uses the "sslget" utility to communicate with the CA. Due to a change in Network Security Services (NSS), the server sent out a full response to the sslget client consisting of 9906 bytes but the client received only 5 bytes of the encrypted stream. With this update the problem is fixed and sslget now prints the returned XML form from the PKI CA as expected.

BZ#806046

Tomcat has changed the way the server startup is logged. In previous versions, server startup and operation was written to the catalina.out file by the root and tomcat users. Now, the root and tomcat users write to different logs. After the change, the Certificate System (CS) tomcat subsystems failed to start due to incorrect permissions. The CS startup code has been modified to reflect this new logging and now works as expected.

All users of pki-core are advised to upgrade to these updated packages, which fix these bugs.

5.246. PM-UTILS

5.246.1. RHBA-2012:1094 – pm-utils bug fix update

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

The pm-utils packages contain a set of utilities and scripts for tasks related to power management.

Bug Fix

BZ#800630

Prior to this update, the RPM description contained wrong product names. This update removes all wrong information.

All users of pm-utils are advised to upgrade to these updated packages, which fix this bug.

5.247. POLICYCOREUTILS

5.247.1. RHBA-2012:0969 – policycoreutils bug fix update

Updated policycoreutils packages that fix several bugs 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.

These updated policycoreutils packages provide fixes for the following bugs:

BZ#784595

The semanage utility did not produce correct audit messages in the Common Criteria certified environment. This update modifies semanage so that it now sends correct audit events when the user is assigned to or removed from a new role.

This update also modifies behavior of semanage concerning the user's SELinux Multi-Level Security (MLS) and Multi-Category Security (MCS) range. The utility now works with the user's default range of the MLS/MCS security level instead of the lowest.

In addition, the semanage(8) manual page has been corrected to reflect the current semanage functionality.

BZ#751313

Prior to this update, the ppc and ppc64 versions of the policycoreutils package conflicted with each other when installed on the same system. This update fixes this bug; ppc and ppc64 versions of the package can now be installed simultaneously.

BZ#684015

The missing exit(1) function call in the underlying code of the sepolgen-ifgen utility could cause the restorecond daemon to access already freed memory when retrieving user's information. This would cause restorecond to terminate unexpectedly with a segmentation fault. With this update, restorecond has been modified to check the return value of the getpwuid() function to avoid this situation.

BZ#786191

When installing packages on the system in Federal Information Processing Standard (FIPS) mode, parsing errors could occur and installation failed. This was caused by the "/usr/lib64/python2.7/site-packages/sepolgen/yacc.py" parser, which used MD5 checksums that are not supported in FIPS mode. This update modifies the parser to use SHA-256 checksums and installation process is now successful.

BZ#786664

Due to a pam_namespace issue which caused a leak of mount points to the parent namespace, polyinstantiated directories could be seen by users other than the owner of that directory. With this update, the mount points no longer leak to the parent namespace, and users can only see directories they own.

BZ#806736, BZ#807011

When a user or a program ran the "semanage fcontext" command, a traceback error was returned. This was due to a typographical error in the source code of the semanage command. This updates fixes this error, and executing the semanage fcontext command works as expected.

All users of policycoreutils are advised to upgrade to these updated packages, which fix these bugs.

5.248. PORTRESERVE

5.248.1. RHBA-2012:0447 – portreserve bug fix update

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

The portreserve package helps services with well-known ports that lie in the portmap range. It prevents portmap from occupying a real service's port by occupying it itself, until the real service tells it to release the port, generally in the init script.

Bug Fixes

BZ#614924

Prior to this update, the init script for the portreserve daemon did not always return the correct exit code. As a consequence, an incorrect error message was displayed. With this update, the init script is modified to return the correct exit codes, and also appropriate messages are now displayed.

BZ#712362

The portreserve package requires the "chkconfig" command because it is run in installation scriptlets. However, this was previously not reflected in the package metadata, and error messages could be displayed during installation. To prevent this issue, this update adds requirement tags for chkconfig.

All users of portreserve are advised to upgrade to this updated package, which fixes these bugs.

5.249. POSTGRESQL AND POSTGRESQL84

5.249.1. RHSA-2012:1037 – Moderate: postgresql and postgresql84 security update

Updated postgresql84 and postgresql packages that fix two security issues are now available for Red Hat Enterprise Linux 5 and 6 respectively.

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

PostgreSQL is an advanced object-relational database management system (DBMS).

Security Fixes

CVE-2012-2143

A flaw was found in the way the crypt() password hashing function from the optional PostgreSQL pgcrypto contrib module performed password transformation when used with the DES algorithm. If the password string to be hashed contained the 0x80 byte value, the remainder of the string was

ignored when calculating the hash, significantly reducing the password strength. This made brute-force guessing more efficient as the whole password was not required to gain access to protected resources.

CVE-2012-2655

Note: With this update, the rest of the string is properly included in the DES hash; therefore, any previously stored password values that are affected by this issue will no longer match. In such cases, it will be necessary for those stored password hashes to be updated.

A denial of service flaw was found in the way the PostgreSQL server performed a user privileges check when applying SECURITY DEFINER or SET attributes to a procedural language's (such as PL/Perl or PL/Python) call handler function. A non-superuser database owner could use this flaw to cause the PostgreSQL server to crash due to infinite recursion.

Upstream acknowledges Rubin Xu and Joseph Bonneau as the original reporters of the CVE-2012-2143 issue.

These updated packages upgrade PostgreSQL to version 8.4.12, which fixes these issues as well as several non-security issues. Refer to the PostgreSQL Release Notes for a full list of changes:

http://www.postgresql.org/docs/8.4/static/release.html

All PostgreSQL users are advised to upgrade to these updated packages, which correct these issues. If the postgresql service is running, it will be automatically restarted after installing this update.

5.249.2. RHSA-2012:1263 – Moderate: postgresql and postgresql84 security update

Updated postgresql84 and postgresql packages that fix two security issues are now available for Red Hat Enterprise Linux 5 and 6 respectively.

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

PostgreSQL is an advanced object-relational database management system (DBMS).

Security Fixes

CVE-2012-3488

It was found that the optional PostgreSQL xml2 contrib module allowed local files and remote URLs to be read and written to with the privileges of the database server when parsing Extensible Stylesheet Language Transformations (XSLT). An unprivileged database user could use this flaw to read and write to local files (such as the database's configuration files) and remote URLs they would otherwise not have access to by issuing a specially-crafted SQL query.

CVE-2012-3489

It was found that the "xml" data type allowed local files and remote URLs to be read with the privileges of the database server to resolve DTD and entity references in the provided XML. An unprivileged database user could use this flaw to read local files they would otherwise not have access to by issuing a specially-crafted SQL query. Note that the full contents of the files were not returned, but portions could be displayed to the user via error messages.

Red Hat would like to thank the PostgreSQL project for reporting these issues. Upstream acknowledges Peter Eisentraut as the original reporter of CVE-2012-3488, and Noah Misch as the original reporter of CVE-2012-3489.

These updated packages upgrade PostgreSQL to version 8.4.13. Refer to the PostgreSQL Release Notes for a list of changes:

All PostgreSQL users are advised to upgrade to these updated packages, which correct these issues. If the postgresql service is running, it will be automatically restarted after installing this update.

5.250. POSTGRESQL-JDBC

5.250.1. RHEA-2012:0991 – postgresql-jdbc enhancement update

An updated postgresql-jdbc package that adds various enhancements is now available.

PostgreSQL is an advanced Object-Relational database management system. The postgresql-jdbc package includes the .jar files needed for Java programs to access a PostgreSQL database.

Enhancement

BZ#816731

This update enables support for the JDBC 4 and JDBC 4.1 extensions of the JDBC API specification. Among other advantages, these additions are necessary for proper functioning of the PostgreSQL JDBC driver under JDK version 6 and later (including JDK 7). This update also converts the driver from a GCJ build to a pure jar (noarch) build.

Users of postgresql-jdbc are advised to upgrade to this updated package, which adds these enhancements.

5.251. PPC64-UTILS

5.251.1. RHEA-2012:0815 – ppc64-utils enhancement update

An updated ppc64-utils package that adds one enhancement is now available for Red Hat Enterprise Linux 6.

The ppc64-utils package is a collection of utilities for Linux running on 64-bit PowerPC platforms.

Enhancement

BZ#739123

By adding the ppc64-diag package into Red Hat Enterprise Linux 6, the ppc64-utils package became dependent on ppc64-diag. This update modifies the spec file of ppc64-utils so that ppc64-diag is now included as a prerequisite of ppc64-utils.

All users of ppc64-utils are advised to upgrade to this updated package, which adds this enhancement.

5.252. PROCPS

5.252.1. RHBA-2012:1463 – procps bug fix update

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

The procps packages provide a set of system utilities to provide system information using the /proc file system. The procps package includes the free, pgrep, pkill, pmap, ps, pwdx, skill, slabtop, snice, sysctl, tload, top, uptime, vmstat, w, and watch utilities.

Bug Fixes

BZ#851664

Prior to this update, the 'si' and 'so' values were always zero for "m" or "M" units. This was caused by an arithmetic precision loss in the expressions used for the calculations. This update modifies the expressions to avoid precision losses.

BZ#875077

Prior to this update, the vmstat tool could be terminated unexpectedly raising the SIGFPE exception when the total sum of 'us', 'sy', 'id', 'wa' and 'st' values returned by the kernel was zero. This situation could only appear on certain specific platforms. this update modifies the internal evaluation so that the vmstat tool is more robust and does no longer terminate.

All users of procps are advised to upgrade to these updated packages, which fix these bugs.

5.252.2. RHBA-2012:0461 – procps bug fix update

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

The procps package contains a set of system utilities that provide system information using the /proc file system. The procps package includes the free, pgrep, pkill, pmap, ps, pwdx, skill, slabtop, snice, sysctl, tload, top, uptime, vmstat, w, and watch utilities.

Bug Fixes

BZ#746997

Prior to this update, it was not possible to instruct the "top" utility by using command-line options to sort processes according to the memory consumption. This was possible only in interactive mode (by using the "Shift+M" key combination). This update introduces a new command-line option, "-a", that instructs the "top" utility to sort processes according to the memory consumption in batch mode. Note that in Red Hat Enterprise Linux 5, this functionality is provided by the "-m" option whereas the same option is used for another feature in Red Hat Enterprise Linux 6.

BZ#751475

Previously, the CPULOOP variable effected only the statistical lines representing the particular CPU cores. With this update, the "top" command additionally applies the CPULOOP variable on the CPU summary line (that represents all the CPU cores).

BZ#766792

Prior to this update, no development package was generated for procps. With this update, a separate procps-devel package containing a set of procps development headers is available.

All users of procps are advised to upgrade to this updated package, which fixes these bugs.

5.253. PSACCT

5.253.1. RHBA-2012:1082 – psacct bug fix update

Updated psacct packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The psacct packages contain utilities for monitoring process activities, including ac, lastcomm, accton, dump-acct, dump-utmp and sa. The "ac" command displays statistics about how long users have been logged on. The "lastcomm" command displays information about previously executed commands. The "accton" command turns process accounting on or off. The "dump-acct" command transforms the output file from the accton format to a human-readable format. The "dump-utmp" command prints utmp files in human-readable format. The "sa" command summarizes information about previously executed commands.

Bug Fixes

BZ#828728

Previously, improper data type detection could have caused an arithmetic overflow. As a consequence, the dump-acct tool reported incorrect elapsed time values. A patch has been applied so that correct values are reported with this update.

BZ#834217

Previously, improper data type conversion caused the dump-utmp tool to report invalid timestamps. Consequently, mainly on the 64-bit PowerPC architecture, dump-utmp could have terminated unexpectedly with a segmentation fault. A patch has been applied so that correct values are reported and no crashes occur with this update.

BZ#838998

Previously, accessing an incorrect and uninitialized memory structure used for acquiring the user IDs caused the sa tool to report incorrect usernames. Consequently, the sa tool could have terminated unexpectedly with a segmentation fault. A patch has been applied so that correct values are reported and no crashes occur with this update.

All users of psacct are advised to upgrade to these updated packages, which fix these bugs.

5.254. PULSEAUDIO

5.254.1. RHBA-2012:1070 – pulseaudio bug fix update

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

PulseAudio is a sound server for Linux and other Unix like operating systems.

Bug Fix

BZ#836139

On certain sound card models by Creative Labs, the S/PDIF Optical Raw output was enabled on boot regardless of the previous settings. This caused the audio output on the analog duplex output to be disabled. With this update, the S/PDIF Optical Raw output is disabled on boot so that the analog output works as expected.

All users of pulseaudio are advised to upgrade to these updated packages, which fix this bug.

5.255. PYKICKSTART

5.255.1. RHBA-2012:0882 – pykickstart bug fix and enhancement update

Updated pykickstart packages that fix multiple bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The pykickstart package contains a python library for manipulating kickstart files.

Bug Fixes

BZ#758603

If using the "raid" command with the "--useexisting" option without specifying the members that were part of the RAID set, the system installation failed with the following error message on system startup:

> Partitions required for raid

With this update, the library scripts check if the RAID members are defined and, if this is not the case, the scripts raise an error.

BZ

It was not possible to configure and use multipath for iSCSI devices using different network interfaces on one subnet during installation as all devices used the default network interface. The "--iface[number]" option has been provided for the "iscsi" command to allow explicit specification of interface binding.

Enhancements

BZ#821315

The "part" command can now take the "--hibernation" option, which allows specifying the size of the swap partition. The existing "--recommended" option follows the Installation Guide swap recommendation, which is not necessarily correct for all systems. This new option allows for a different sizing algorithm that can be more suitable, especially if the system uses hibernation.

BZ#790457

Logical Volume groups can now have free space reserved using either the "--reserved-space" or "--reserved-percent" parameters. The "--reserved-space" option takes a number representing megabytes as its argument, while the "--reserved-percent" option takes a number representing the percentage of the volume group that should be left free. The options can only be used for volume groups created during installation.

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

5.256. PYQT4

5.256.1. RHBA-2012:1241 – PyQt4 bug fix update

Updated PyQt4 packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The PyQt4 packages contain python bindings for Qt4.

Bug Fixes

BZ#757411

Prior to this update, the PyQt4 utility did not contain the deleteResource method of PyQt4.QtNetwork.QNetworkAccessManager. This update modifies the underlying code to include the missing qnetwork-deleteResource method.

BZ#821061

Prior to this update, the PyQt4 utility did not contain the QMenuBar.setCornerWidget method. This update modifies the underlying code to include the missing qmenubar-cornerWidget method.

All users of PyQt4 are advised to upgrade to these updated packages, which fix these bugs.

5.257. PYTHON-CONFIGSHELL

5.257.1. RHBA-2012:0856 – python-configshell bug fix and enhancement update

Updated python-configshell packages that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The python-configshell packages provide a library for implementing configuration command line interfaces for the Python programming environment.

The python-configshell package has been upgraded to version 1.1.fb4 which provides a number of bug fixes and enhancements over the previous version, and adds support for the configuration shell functionality of fcoe-target-utils packages. (BZ#765977)

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

5.258. PYTHON-MEMCACHED

5.258.1. RHBA-2012:0889 – python-memcached bug fix update

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

The python-memcached packages provide a python interface to the memcached memory cache daemon.

Bug Fix

BZ#789494

Prior to this update, python-memcached failed to get statistics for the memcached server and the "get_stats()" function returned a warning message. This update modifies the "get_stats()" method so that python-memcached can now get the statistics as expected.

All users of python-memcached are advised to upgrade to these updated packages, which fix this bug.

5.259. PYTHON-PASTE-SCRIPT

5.259.1. RHSA-2012:1206 – Moderate: python-paste-script security update

An updated python-paste-script package that fixes one security issue is now available for Red Hat Enterprise Linux 6.

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

Python Paste provides middleware for building and running Python web applications. The python-paste-script package includes paster, a tool for working with and running Python Paste applications.

Security Fix

CVE-2012-0878

It was discovered that paster did not drop supplementary group privileges when started by the root user. Running "paster serve" as root to start a Python web application that will run as a non-root user and group resulted in that application running with root group privileges. This could possibly allow a remote attacker to gain access to files that should not be accessible to the application.

All paster users should upgrade to this updated package, which contains a backported patch to resolve this issue. All running paster instances configured to drop privileges must be restarted for this update to take effect.

5.260. PYTHON-REPOZE-WHO

5.260.1. RHEA-2012:0982 – python-repoze-who enhancement update

An updated python-repoze-who package that adds one enhancement is now available for Red Hat Enterprise Linux 6.

The python-repoze-who package provides an identification and authentication framework for arbitrary Python WSGI applications and acts as WSGI middleware.

The python-repoze-who package has been upgraded to upstream version 1.0.18, which adds support for enforcing timeouts for cookies issued by the auth_tkt plugin via the "timeout" and "reissue_time" config parameters. (BZ#639075)

All users of python-repoze-who are advised to upgrade to this updated package, which adds this enhancement.

5.261. PYTHON-RHSM

5.261.1. RHBA-2012:0805 – python-rhsm bug fix and enhancement update

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

The python-rhsm package contains a library for communicating with the representational state transfer (REST) interface of Red Hat's subscription and content service. This interface is used by the

Subscription Management tools for management of system entitlements, certificates, and access to content.

Bug Fixes

BZ#720372

Previously, Subscription Manager had the fakamai-cp1.pem certificate installed in the /etc/rhsm/ca/fakamai-cp1.pem directory after installation. However, the certificate serves only for testing purposes and is not needed by the tool itself. With this update, the certificate has been removed.

BZ#744654

If the subscription-manager command was issued with an incorrect or empty --server.port option, the command failed with a traceback. With this update, the tool now sets the provided port value as expected and no traceback is returned.

BZ#803773

If the activation key contained non-ASCII characters, the registration failed with the following error:

> Network error. Please check the connection details, or see /var/log/rhsm/rhsm.log for more information.

This happened due to an incorrect conversion of the key into the URL address. With this update, subscription-manager converts the characters correctly and the registration succeeds in the scenario described.

BZ#807721

Several configuration settings had no default value defined in the Red Hat Subscription Manager (RHSM), which could cause some commands to return a traceback. The default RHSM values are now set as expected and the problem no longer occurs.

BZ#822965

When the user defined a proxy server in rhsm.conf, the Subscription Manager did not work and returned the "unknown URL type" error. This happened because the "Host" header was not sent up to the CDN when acquiring the list of releases using a proxy. With this update, the "Host" header is sent to the CDN and the proxy definition in rhsm.conf is processed as expected.

Enhancement

BZ#785247

The bug fixes and some of the new features introduced in the python-rhsm package on Red Hat Enterprise Linux 5.8 have been backported into the python-rhsm on Red Hat Enterprise Linux 6.3.

All users of python-rhsm are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.262. PYTHON-RTSLIB

5.262.1. RHBA-2012:0855 – python-rtslib bug fix update

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

The python-rtslib packages provide a Python library to configure the kernel target subsystem with the configfs file system.

Bug Fix

BZ#813676

Prior to this update, configurations were not saved when the targetcli tool exited. This update improves the mechanism for saving and restoring target configuration across system restarts.

All users of python-rtslib are advised to upgrade to these updated packages, which fix this bug.

5.263. PYTHON

5.263.1. RHBA-2012:1250 – python bug fix update

Updated python packages that fix a bug are now available for Red Hat Enterprise Linux 6.

Python is an interpreted, interactive, object-oriented programming language. Python includes modules, classes, exceptions, high-level dynamic data types, and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing systems (X11, Motif, Tk, Mac and MFC).

Bug Fix

BZ#848815

As part of the fix for CVE-2012-0876, a new symbol ("XML_SetHashSalt") was added to the system libexpat library, and which Python's standard library uses within the pyexpat module. If an unpatched libexpat.so.1 was present in a directory listed in LD_LIBRARY_PATH, then attempts to use the pyexpat module (such as within yum) would fail with an ImportError exception. This update adds an RPATH directive to pyexpat to ensure that the system libexpat is used by pyexpat, regardless of whether there is an unpatched libexpat within the LD_LIBRARY_PATH, thus preventing the ImportError exception.

All Python users are advised to upgrade to these updated packages, which fix this bug.

5.264. PYTHON-VIRTINST

5.264.1. RHBA-2012:0784 – python-virtinst bug fix update

Updated python-virtinst packages that fix several bugs 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#741158

Previously, some special characters could not be included in virtual machine descriptions in virt-manager (the Virtual Machine Manager graphical tool for administration of virtual machines). Saving a description containing such characters failed and the following error message was reported:

-

> Error changing VM configuration: 'NoneType' object is not callable

This update fixes the error so that descriptions containing special characters can be saved correctly and the error message no longer occurrs.

BZ#765928

Previously, creation of guests in a multi-threaded application like virt-manager sometimes failed with the following error message:

> RuntimeError: dictionary changed size during iteration

This was caused by improper locking of internal virtinst state. This update fixes the error and as a result, the aforementioned error no longer occurs when creating guests in multi-threaded applications.

BZ#769191

Previously, incorrect IDs were generated by virt-manager for newly created virtio hard disks when 26 or more virtio hard disks were already created. Consequently, creation of the 27th virtio hard disk failed with the following error message:

> Error adding device: An error occurred, but the cause is unknown

This update ensures that virt-manager generates correct IDs for all virtio hard disks. As a result, creation of virtio hard disks is possible up to the supported virtio hard disk limit.

BZ#783866

Prior to this update, Japanese translation of the --help message for virt-clone was incomplete and a part of it was displayed in English. Also, the meter message for cloning progress in virt-clone displayed unknown characters when viewed in a terminal window of a certain size. This update fixes both problems and ensures that both messages are now translated and displayed correctly.

BZ#786672

Prior to this update, virt-manager expected a string value when returning virtual machine descriptions. However, virtual machines created with virt-manager-0.9.0-7.el6 or older did not have any description, and therefore returned a "None" value instead. Consequently, loading of guests created with virt-manager-0.9.0-7.el6 or older using a newer version failed with the following error message:

> AttributeError: 'NoneType' object has no attribute 'replace'

This update adds a return value check to the operation of loading defined guests and ensures backward compatibility of guests created with virt-manager-0.9.0-7.el6 or older. As a result, these guests are now loaded correctly.

BZ#798909

Prior to this update, virt-clone recognized only the "xen" and "qemu" domain types, while "xen" was the default. Consequently, when "kvm" was specified as a domain type of a cloned virtual machine, virt-clone did not recognize the domain type and used "xen" instead. This update adds support for the "kvm" domain type. As a result, "kvm" is now recognized when specified as a domain type.

All users of python-virtinst are advised to upgrade to these updated packages, which fix these bugs.

5.265. QEMU-KVM

5.265.1. RHSA-2012:1234 – Important: qemu-kvm security update

Updated qemu-kvm packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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(s) associated with each description below.

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.

Security Fix

CVE-2012-3515

A flaw was found in the way QEMU handled VT100 terminal escape sequences when emulating certain character devices. A guest user with privileges to write to a character device that is emulated on the host using a virtual console back-end could use this flaw to crash the qemu-kvm process on the host or, possibly, escalate their privileges on the host.

This flaw did not affect the default use of KVM. Affected configurations were:

- When guests were started from the command line ("/usr/libexec/qemu-kvm") without the "-nodefaults" option, and also without specifying a serial or parallel device, or a virtio-console device, that specifically does not use a virtual console (vc) back-end. (Note that Red Hat does not support invoking "qemu-kvm" from the command line without "-nodefaults" on Red Hat Enterprise Linux 6.)

- Guests that were managed via libvirt, such as when using Virtual Machine Manager (virt-manager), but that have a serial or parallel device, or a virtio-console device, that uses a virtual console back-end. By default, guests managed via libvirt will not use a virtual console back-end for such devices.

Red Hat would like to thank the Xen project for reporting this issue.

All users of qemu-kvm should upgrade to these updated packages, which resolve this issue. 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.

5.265.2. RHBA-2012:1585 – qemu-kvm bug fix update

Updated qemu-kvm packages that fix two 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. The qemu-kvm packages form the user-space component for running virtual machines using KVM.

Bug Fixes

BZ#861049

If no listener is connected to a port on a host, output from the guest is suppressed until a listener is connected. However, for console ports, the guest output needs to be discarded instead. Previously, the guest kept waiting for a listener after it wrote data to a console port. But since there was no listener, the guest eventually became unresponsive. This bug has been fixed by changing behavior of the "pty" socket type to not suppress output from the ports and properly discard the data if no listener is connected.

BZ#861906

With some initial guest OS installations using the QXL driver and VNC as the display protocol, virtual machines were terminating unexpectedly with a segmentation fault during setup and returned the "lost connection with kvm process" error message. A patch has been provided to address this issue and virtual machines now run properly in the described scenario.

All users of qemu-kvm are advised to upgrade to these updated packages, which fix these bugs. 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.

5.265.3. RHBA-2012:1519 – qemu-kvm bug fix update

Updated qemu-kvm packages that fix a bug 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. The qemu-kvm packages form the user-space component for running virtual machines using KVM.

Bug Fix

BZ#873270

In the SVVP (Server Virtualization Validation Program) environment, when the e1000 network driver was used, the PCI Hardware Compliance Test For Systems job failed. Consequently, the HCK (Hardware Certification Kit) SVVP certification could not be passed on the system. A patch has been provided to address this issue and the test now passes in the described scenario.

All users of qemu-kvm are advised to upgrade to these updated packages, which fix this bug. 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.

5.265.4. RHBA-2012:1582 – qemu-kvm bug fix update

Updated qemu-kvm packages that fix a bug 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. The qemu-kvm packages form the user-space component for running virtual machines using KVM.

Bug Fix

BZ#886101

When the vdsm daemon was running on a blocking NFS storage, it attempted to continuously access the storage. Consequently, vdsm could become unresponsive for almost an hour. This bug has been fixed and vdsm is now able to recover within a few minutes in the described scenario.

All users of qemu-kvm are advised to upgrade to these updated packages, which fix this bug. 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.

5.265.5. RHBA-2012:1121 – qemu-kvm bug fix update

Updated qemu-kvm packages that fix two 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. The qemu-kvm packages form the user-space component for running virtual machines using KVM.

Bug Fixes

BZ#839897

Previously, the KVM modules were not loaded by the postinstall scriplet of RPM scripts. This bug caused various issues and required the system to be rebooted to resolve them. With this update, the modules are loaded properly by the scriptlet and no unnecessary reboots are now required.

BZ#840054

Previously, when a guest was started up with two serial devices, qemu-kvm returned an error message and terminated the boot because IRQ 4 for the ISA bus was being used by both devices. This update fixes the qemu-kvm code, which allows IRQ 4 to be used by more than one device on the ISA bus, and the boot now succeeds in the described scenario.

All users of qemu-kvm are advised to upgrade to these updated packages, which fix these bugs. 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.

5.265.6. RHBA-2012:0746 – qemu-kvm bug fix and enhancement update

Updated qemu-kvm packages that fix multiple bugs and add various enhancements 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. The qemu-kvm packages form the user-space component for running virtual machines using KVM.

Bug Fixes

BZ#787974

A virtio-serial device marked a guest driver with the "present" bit even if a guest was not present. Because the bit was set, the Simple Protocol for Independent Computing Environments (SPICE) assumed, after the migration process had completed, that a guest agent was running, and disabled the server-side mouse. This caused the mouse to be unusable even if no agent was running. The "present" bit is now only set if a working guest driver is present. When no guest agent is running, the mouse continues working after migration as expected.

BZ#789417

Previously, the free() function was missing in management of the "xsave" processor state. This led to memory leaks in qemu-kvm when a guest used the xsave functionality, causing excessive memory consumption on the host. Buffers used to manage xsave support are now freed after use so that

qemu-kvm no longer leaks memory.

BZ#790421

It is possible to set up SPICE channels using Transport Layer Security (TLS) when no TLS port has been specified (that is, TLS is disabled). Due to this, it was previously not possible to connect to a virtual machine using SPICE when starting QEMU. With this update, QEMU now exits with an error message in such a situation.

BZ#807313

Previously, in certain cases, the USB storage emulation feature failed to update the state correctly on I/O request cancellation. As a consequence, the USB storage machine triggered an assertion in the USB core code, leading to the qemu process dumping core. With this update, status updates are handled correctly, and qemu no longer dumps core.

BZ#807916

The QEMU Enhanced Host Controller Interface (EHCI) code previously contained a spurious assert() function. As a consequence, qemu could dump core. The assert() function has been removed, preventing core dumps in this scenario.

BZ#748810

Due to a bug in the QXL driver, if the user started a QEMU guest, stopped the guest, and executed the "screendump" command, the qemu-kvm process terminated unexpectedly with a segmentation fault. The bug in the QXL driver has been fixed, and qemu-kvm no longer crashed in this scenario.

BZ#785963

If the user pressed a modifier key (Shift, Ctrl or Alt) while closing a Virtual Network Computing (VNC) connection, the key event was treated as if pressed when the next VNC connection was opened. This happened for example when the VNC viewer was closed with the Alt+F4 key combination. To prevent this problem, a key-up event is now injected into the guest and the event is handled as expected if any modifier key is pressed when closing the VNC connection.

BZ#738519

When hot plugging or hot unplugging a USB controller more than 1000 times, the qemu-kvm process dumped core. This was because the Memory-mapped I/O (MMIO) BARs were present, but failed to be unregistered. Unregistering of MMIO BARs has been implemented with this update, so that qemu-kvm runs correctly and a USB controller can be hot plugged and hot unplugged multiple times as expected.

BZ#740707

Due to an assertion performed on packet completion, running a guest with a USB device passthrough to the USB 1.1 controller caused qemu-kvm to terminate with an assertion failure. The assertion is no longer performed on packet completion, which ensures that qemu-kvm runs correctly.

BZ#734426

When starting a guest and moving time backwards on the host, the guest became unresponsive. This was because of incorrect real-time clock (RTC) timer emulation. This problem has been fixed, and the guest settings are now adjusted properly so the guest does not hang in this scenario.

BZ#795652

If the "__com.redhat_spice_migrate_info" monitor command was applied with incorrect parameters,

the error handler caused the QEMU monitor to become unresponsive. Error handling has been modified so that the QEMU monitor no longer hangs when executing commands with incorrect parameters.

BZ#796063

An incorrect bit was set in the SAVE/RESTORE handler. As a consequence, the guest could become unresponsive after a live or save and restore migration. A patch has been applied to address this issue, so that the guest no longer hangs in this scenario.

BZ#754349

USB device initialization failure was not handled properly. Adding multiple invalid USB host devices led to the guest dumping core. USB initialization failure handling has been fixed so that the guest no longer dumps core under these circumstances.

BZ#702370

Due to incorrect calculation of transferred bytes, migration downtime was previously longer than expected and allowed by setting the "migrate_max_downtime()" monitor command. A guest was therefore unavailable for much longer time than allowed. The underlying source code has been modified to calculate only transferred bytes, so that qemu-kvm now honors migration downtime settings.

BZ#698936

Migration to hosts with earlier versions of Red Hat Enterprise Linux (notably Red Hat Enterprise Linux 6.1) could fail due to incompatible QXL revision. The revision number has been changed to be compatible with older versions of Red Hat Enterprise Linux, so that guests can now be successfully migrated to such hosts.

BZ#769760

The USB controller did not wait to finish the last transaction before trying to format a USB device. If any USB operation was in progress when resetting the USB device, the operation never finished. This update fixes detaching of a child process, so that the reset process starts after the transaction has finished. Formatting a usb-storage device and the respective USB operations therefore finish successfully.

BZ#769745

Previously, the USB release function was not called in the exit notifier. As a consequence, the host was unable to reuse the USB device after it had been removed from the guest. With this update, the release function is now called in the exit notifier, ensuring that the host can reuse the USB after it has been removed from the guest.

BZ#796118

Previously, the QEMU USB emulation code modified data structures after releasing them. Consequently, an assertion was triggered due to unexpected data structure changes, and the qemu process dumped core. The release call has been moved to the correct place in the code, so that core dumps no longer occur in this scenario.

BZ#769142

When using VNC reverse mode (QEMU connects to the VNC viewer, not vice versa), the VNC server attempted to access the display before initialization. This led to a core dump on the guest machine. With this update, the display is initialized before it is used.

BZ#638055

In safe mode, the "qemu-img rebase" command incorrectly handled backing files as if they were of the same size as the rebased image. As a consequence, attempting to rebase an image if the old or the new backing file was smaller than the image itself failed with the following error message:

```
qemu-img: error while reading from new backing file
```

With this update, backing files are handled correctly, and the "qemu-img rebase" command succeeds even if a backing file is smaller than the rebased image.

BZ#736942

The cleanup code of the qemu-img utility did not perform NULL pointer checks for old and new backing files. When executing "qemu-img rebase" in safe mode on an image with a backing file that could not be opened, the utility terminated unexpectedly with a segmentation fault and an error message. This update adds the necessary NULL pointer checks to the cleanup code of qemu-img, and qemu-img now exits gracefully if either the new or the old backing file cannot be opened.

BZ#737879

Due to incorrect handling of invalid arguments for the "-drive" option, running "qemu-kvm -drive" with such arguments could lead to a drive misconfiguration. Validation of the "-drive" arguments has been corrected. As a result, qemu-kvm fails to run if invalid arguments are used.

BZ#790083

When migrating a Microsoft Windows guest with the QXL display driver version 0.1.9 or earlier, the screen of the destination machine could contain rendering artifacts because the primary surface memory was not up to date. With this update, the primary surface memory is updated properly, and screen corruption no longer occurs under these circumstances.

BZ#781920

While reallocating transmission buffers, the guest driver could have allocated unlimited memory to the transmission buffers. This caused qemu to terminate with a glib error. The transmission buffer size is now limited, which prevents the guest from allocating unlimited memory to the buffers, and qemu no longer crashes under these circumstances.

BZ#796575

Previously, the qemu process required periodic polling for events, which could lead to qemu waking up multiple times per a second. Because qemu can set specific timers to poll for events periodically, the generic polling timer has been removed. As a result, an idle guest with no VNC or SPICE connections active does not wake up the qemu process unnecessarily.

BZ#798936

Running the "qemu -cpu host" command did not expose emulated Performance Monitor Unit (PMU) to a guest, and the guest was therefore unable to use the PMU counters. With this update, if the host kernel supports PMU emulation, the CPUID OAH leaf is exposed to a guest, so the guest can use the PMU counters to profile itself.

BZ#757713

The code for monitor file name completion previously incorrectly checked for a directory. Consequently, the slash character could be appended to a string even if the completed name did not refer to a directory. The check for directory has been fixed, and the slash character is now added only when the completed name refers to a directory.

BZ#757132

Implementation of the VGA underline attribute could read beyond arrays and corrupted pixels in underlined characters could have been observed on a guest running a non-framebuffer text console. With this update, reading beyond arrays is no longer allowed, and corrupted pixels are no longer present.

BZ#752049

The Enhanced Host Controller Interface (EHCI) reset handler was incorrect. Microsoft Windows guests could become unresponsive on boot with USB disk passthrough when loading the USB controller. Reset handler implementation has been fixed, and Microsoft Windows guests with USB device passthrough now boot successfully.

BZ#749820

A use-after-free bug in the "acl_reset" monitor command could cause the qemu-kvm process to terminate unexpectedly with a segmentation fault. With this update, the use of freed memory is avoided, and qemu-kvm no longer crashes under these circumstances.

BZ#747010

An incorrect value was used to calculate memory usage of the qemu-kvm process and to turn on Kernel Samepage Merging (KSM). As a consequence, KSM was turned on too early. Real memory size is now used instead of virtual memory size for calculating qemu-kvm memory usage, which ensures that turning on KSM is now more optimized.

BZ#743251

When running the "qemu-kvm" command without the "-spice" option, qemu-kvm terminated unexpectedly with a segmentation fault if the user attempted to run the "info spice" monitor command afterward. A check has been added to verify whether the command is being run with the "-spice" option, so that qemu-kvm no longer crashes in this scenario.

BZ#812328

Enhanced Host Controller Interface (EHCI) emulation previously had a limitation of the number of queue heads processed. If many devices were present, EHCI did not process all queues, rendering some devices non-functional. The limitation has been removed, so that EHCI now works as expected with a high number of devices.

BZ#728385

When running qemu-kvm with the "-nographic" option, and then executing the "screendump" command, qemu-kvm terminated unexpectedly with a segmentation fault. A check for a valid screendump function pointer has been implemented, and is performed prior to calling the function. As a result, qemu-kvm no longer crashes in this scenario.

Enhancements

BZ#562886

KVM now supports dynamic CPU allocation, also called vCPU hot plug, as a Technology Preview. This feature allows users to dynamically adjust CPU resources in a guest. The availability of the guest is increased, because it is no longer needed to take the guest offline to adjust CPU resources.

BZ#632771

Now, qemu-kvm contains a new sub-package called qemu-guest-agent. When running Red Hat

Enterprise Linux 6.3 guests with this package installed, properly configured Red Hat Enterprise Linux 6.3 hosts can send new commands to the guest, for example, "guest-sync", "guest-ping", "guest-info", "guest-shutdown", and "guest-suspend-*".

BZ#783950

KVM in Red Hat Enterprise Linux 6.3 now has an improved access to qcow2 disk images (qcow2 is the default format), which is now more asynchronous. The vCPU stalls frequency has been thus decreased resulting in an overall performance improvement during disk I/O.

BZ#782029

KVM Virtualization's storage stack has been improved with the addition of virtio-SCSI (a storage architecture for KVM based on SCSI) capabilities. Virtio-SCSI now provides the ability to connect directly to SCSI LUNs and significantly improves scalability compared to virtio-blk. The advantage of virtio-SCSI is that it is capable of handling hundreds of devices compared to virtio-blk which can only handle 28 devices and exhausts PCI slots.

BZ#767302

This update adds new CPU model definitions for the latest AMD processors.

BZ#760953

This update adds new CPU model definitions for Intel Core i3, i5 and i7 processors.

BZ#758104

Spice builds on KVM USB 2.0 host adapter emulation support, and enables remote USB redirection support that allows virtual machines running on servers to use remotely plugged USB devices on the client side.

All users of qemu-kvm are advised to upgrade to these packages, which fix these bugs and add these enhancements.

5.266. QL2400-FIRMWARE

5.266.1. RHBA-2012:0859 – ql2400-firmware bug fix and enhancement update

An updated ql2400-firmware package that provides several bug fixes and 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.

This update upgrades the ql2400 firmware to upstream version 5.06.05, which provides a number of bug fixes and enhancements over the previous version. (BZ#766048)

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.

5.267. QL2500-FIRMWARE

5.267.1. RHBA-2012:0860 – ql2500-firmware bug fix update

An updated ql2500-firmware package that fixes multiple 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.

This update upgrades the ql2500 firmware to upstream version 5.06.05, which provides a number of bug fixes and enhancements over the previous version. (BZ#766050)

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.

5.268. QPID-CPP, PYTHON-QPID, AND SASLWRAPPER

5.268.1. RHBA-2012:0764 – qpid-cpp, python-qpid and saslwrapper bug fix and enhancement update

Updated qpid-cpp, python-qpid, and saslwrapper packages that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

Qpid provides a message broker daemon that receives, stores and routes messages using the open AMQP messaging protocol along with run-time libraries for AMQP client applications developed using Qpid C++. Clients exchange messages with an AMQP message broker using the AMQP protocol.

All qpid packages have been upgraded to upstream version 0.14, which provides a number of bug fixes and enhancements over the previous version. (BZ#765803, BZ#765863)

Bug Fixes

BZ#785919

Prior to this update, the spec file incorrectly used the "Vendor" tag. As a consequence, rebuilds could, under certain circumstances, incorrectly label packages. This update removes the "Vendor" tag.

BZ#788901

Prior to this update, the qpid daemon could, under certain circumstances, abort with a segmentation fault when attempting to shut down the SslPlugin if the plug-in was loaded without setting the "--ssl-cert-db" option. This update modifies the ssl.so module so that the SslPlugin shuts now down as expected.

BZ#799269

Prior to this update, the service qpid daemon was always enabled by default, which was not required. This update modifies the service settings so that the qpidd service is no longer enabled by default but only when required.

Enhancements

BZ#703563

This update provides the environment variable "QPIDC_CONF_FILE" to allow to specify the location to look for the configuration file of the client.

BZ#749600

This update adds the python-saslwrapper package as a dependency for qpid-stat so that qpid-stat can use the DIGEST-MD5 module to authenticate to a broker.

BZ#771961

With this update, the spec file depends on specific boost-* packages instead of the boost metapackage.

BZ#808783

This update adds support for 64-bit PowerPC and IBM System z platforms to the python-saslwrapper package.

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

5.269. QPID

5.269.1. RHSA-2012:1269 – Moderate: qpid security, bug fix, and enhancement update

Updated qpid packages that fix one security issue, multiple 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

Apache Qpid is a reliable, cross-platform, asynchronous messaging system that supports the Advanced Message Queuing Protocol (AMQP) in several common programming languages.

Security Fix

CVE-2012-2145

It was discovered that the Qpid daemon (qpidd) did not allow the number of connections from clients to be restricted. A malicious client could use this flaw to open an excessive amount of connections, preventing other legitimate clients from establishing a connection to qpidd.

To address CVE-2012-2145, new qpidd configuration options were introduced: max-negotiate-time defines the time during which initial protocol negotiation must succeed, connection-limit-per-user and connection-limit-per-ip can be used to limit the number of connections per user and client host IP. Refer to the qpidd manual page for additional details.

In addition, the qpid-cpp, qpid-qmf, qpid-tools, and python-qpid packages have been upgraded to upstream version 0.14, which provides support for Red Hat Enterprise MRG 2.2, as well as a number of bug fixes and enhancements over the previous version. (BZ#840053, BZ#840055, BZ#840056, BZ#840058)

All users of qpid are advised to upgrade to these updated packages, which fix these issues and add these enhancements.

5.270. QT

5.270.1. RHBA-2012:1246 – qt bug fix update

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

The qt packages contain a software toolkit that simplifies the task of writing and maintaining GUI (Graphical User Interface) applications for the X Window System.

Bug Fixes

BZ#678604

Prior to this update, the mouse pointer could, under certain circumstances, disappear when using the IRC client Konversation. This update modifies the underlying codes to reset the cursor on the parent and set the cursor on the new window handle. Now, the mouse pointer no longer disappears.

BZ#847866

Prior to this update, the high precision coordinates of the QTabletEvent class failed to handle multiple Wacom devices. As a consequence, only the device that was loaded first worked correctly. This update modifies the underlying code so that multiple Wacom devices are handled as expected.

All users of qt are advised to upgrade to these updated packages, which fix this bugs.

5.270.2. RHSA-2012:0880 – Moderate: qt security and bug fix update

Updated qt packages that fix two security issues and three 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

Qt is a software toolkit that simplifies the task of writing and maintaining GUI (Graphical User Interface) applications for the X Window System. HarfBuzz is an OpenType text shaping engine.

Security Fixes

CVE-2011-3922

A buffer overflow flaw was found in the harfbuzz module in Qt. If a user loaded a specially-crafted font file with an application linked against Qt, it could cause the application to crash or, possibly, execute arbitrary code with the privileges of the user running the application.

CVE-2010-5076

A flaw was found in the way Qt handled X.509 certificates with IP address wildcards. An attacker able to obtain a certificate with a Common Name containing an IP wildcard could possibly use this flaw to impersonate an SSL server to client applications that are using Qt. This update also introduces more strict handling for hostname wildcard certificates by disallowing the wildcard character to match more than one hostname component.

Bug Fixes

BZ#694684

The Phonon API allowed premature freeing of the media object. Consequently, GStreamer could terminate unexpectedly as it failed to access the released media object. This update modifies the underlying Phonon API code and the problem no longer occurs.

BZ#757793

Previously, Qt could output the "Unrecognized OpenGL version" error and fall back to OpenGL-version-1 compatibility mode. This happened because Qt failed to recognize the version of OpenGL installed on the system if the system was using a version of OpenGL released later than the Qt version in use. This update adds the code for recognition of OpenGL versions to Qt and if the OpenGL version is unknown, Qt assumes that the last-known version of OpenGL is available.

BZ#734444

Previously Qt included a compiled-in list of trusted CA (Certificate Authority) certificates, that could have been used if Qt failed to open a system's ca-bundle.crt file. With this update, Qt no longer includes compiled-in CA certificates and only uses the system bundle.

Users of Qt should upgrade to these updated packages, which contain backported patches to correct these issues. All running applications linked against Qt libraries must be restarted for this update to take effect.

5.271. QUAGGA

5.271.1. RHSA-2012:1259 – Moderate: quagga security update

Updated quagga packages that fix multiple security issues 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

Quagga is a TCP/IP based routing software suite. The Quagga bgpd daemon implements the BGP (Border Gateway Protocol) routing protocol. The Quagga ospfd and ospf6d daemons implement the OSPF (Open Shortest Path First) routing protocol.

Security Fixes

CVE-2011-3327

A heap-based buffer overflow flaw was found in the way the bgpd daemon processed malformed Extended Communities path attributes. An attacker could send a specially-crafted BGP message, causing bgpd on a target system to crash or, possibly, execute arbitrary code with the privileges of the user running bgpd. The UPDATE message would have to arrive from an explicitly configured BGP peer, but could have originated elsewhere in the BGP network.

CVE-2011-3323

A stack-based buffer overflow flaw was found in the way the ospf6d daemon processed malformed Link State Update packets. An OSPF router could use this flaw to crash ospf6d on an adjacent router.

CVE-2011-3324

A flaw was found in the way the ospf6d daemon processed malformed link state advertisements. An OSPF neighbor could use this flaw to crash ospf6d on a target system.

CVE-2011-3325

A flaw was found in the way the ospfd daemon processed malformed Hello packets. An OSPF neighbor could use this flaw to crash ospfd on a target system.

CVE-2011-3326

A flaw was found in the way the ospfd daemon processed malformed link state advertisements. An OSPF router in the autonomous system could use this flaw to crash ospfd on a target system.

CVE-2012-0249

An assertion failure was found in the way the ospfd daemon processed certain Link State Update packets. An OSPF router could use this flaw to cause ospfd on an adjacent router to abort.

CVE-2012-0250

A buffer overflow flaw was found in the way the ospfd daemon processed certain Link State Update packets. An OSPF router could use this flaw to crash ospfd on an adjacent router.

CVE-2012-0255, CVE-2012-1820

Two flaws were found in the way the bgpd daemon processed certain BGP OPEN messages. A configured BGP peer could cause bgpd on a target system to abort via a specially-crafted BGP OPEN message.

Red Hat would like to thank CERT-FI for reporting CVE-2011-3327, CVE-2011-3323, CVE-2011-3324, CVE-2011-3325, and CVE-2011-3326; and the CERT/CC for reporting CVE-2012-0249, CVE-2012-0250, CVE-2012-0255, and CVE-2012-1820. CERT-FI acknowledges Riku Hietamäki, Tuomo Untinen and Jukka Taimisto of the Codenomicon CROSS project as the original reporters of CVE-2011-3327, CVE-2011-3323, CVE-2011-3324, CVE-2011-3325, and CVE-2011-3326. The CERT/CC acknowledges Martin Winter at OpenSourceRouting.org as the original reporter of CVE-2012-0249, CVE-2012-0250, and CVE-2012-0255, and Denis Ovsienko as the original reporter of CVE-2012-1820.

Users of quagga should upgrade to these updated packages, which contain backported patches to correct these issues. After installing the updated packages, the bgpd, ospfd, and ospf6d daemons will be restarted automatically.

5.272. QUOTA

5.272.1. RHBA-2012:1472 – quota bug fix update

Updated quota packages that fix multiple bugs 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 Fixes

BZ#680919

Prior to this update, warnquota sent emails from <root@myhost.com> if the quota limit was exceeded and the warnquota tool was enabled to send warning emails and the default warnquota configuration was not changed. As a consequence, users could wrongly reply to this address and email bounces

were delivered to the mailbox of <root@myhost.com>. This update modifies the default warnquota configuration to use the reserved domain "example.com".

BZ#683554

Prior to this update, the option "-r" for setquota and edquota failed to set the grace times for NFS-mounted file systems without reporting errors because the underlying remote procedure call protocol does not support this option. This update disables the option "-r". With this update, the option to set grace times over the network is disabled and error messages are sent when using the "-r" option.

BZ#692390

Prior to this update, the quotacheck tool could mishandle UIDs in processed fsv1 quota files if a user's block limit was reached. This update zeroes uninitialized padding in the "v2r1 ddquot" structure before running subsequent checks.

BZ#704216

Prior to this update, the edquota tool could abort with a segmentation fault if the name server switch was configured to use the libdb back end. This update modifies the underlying code to make the "dirname" symbol in edquota sources static to avoid pollution of the symbol name space confusing the dynamic linker. Now, edquota runs on systems which use the Berkeley DB (BDB) database for storing user names, group names, or passwords.

BZ#730057

Prior to this update, the quota_nld service logged the error message "Failed to find tty of [UID] to report warning to" when users without an interactive session exceeded the disk quota limit while running quota_nld service. This update applies these warnings to non-daemon debugging mode of quota_nld only.

BZ#770307

Prior to this update, the warnquota tool sent a badly worded email message. This update changes the wording and the text is now worded more representative.

All users of quota are advised to upgrade to these updated packages, which fix these bugs.

5.273. RDESKTOP

5.273.1. RHBA-2012:1276 – rdesktop bug fix update

Updated rdesktop packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The rdesktop packages provide a client for the Remote Desktop Server in Microsoft Windows. The rdesktop client uses the Remote Desktop Protocol (RDP) to remotely present a user's desktop.

Bug Fixes

BZ#680917, BZ#772020

Prior to this update, redundant conversions functions did not handle the PC/SC (Personal Computer/Smart Card) integration correctly. As a consequence, the rdesktop on AMD64 and Intel 64 platforms failed to connect and incorrectly. This update removes these redundant functions. This

update also adds smart card reader support for AMD64 and Intel 64 platforms. Now, the rdesktop connects as expected.

BZ#680926

Prior to this update, the rdesktop code for smart card integration with PC/SC caused a buffer overflow on AMD64 and Intel 64 platforms. As a consequence, the glibc function "free()" was aborted with a segmentation fault. This update uses the correct structure and the glibc function "free()" works now as expected.

BZ#782494

Prior to this update, the server generated a cursor-related command that the rdesktop client did not support when using rdesktop to connect to Windows Server 2008 R2 platforms. As a consequence, the mouse pointer was all black. With this update, the mouse pointer is drawn correctly when connecting to Windows Server 2008 R2.

BZ#820008

Prior to this update, the specification file incorrectly listed the libao-devel package as an install dependency for rdesktop. This update removes the libao-devel dependency from the specification file.

BZ#831095

Prior to this update, the rdesktop client did not handle the licenses correctly, As a consequence, certain Terminal Services failed to connect after the first connection with the error message "disconnect: Internal licensing error". This update modifies the underlying code to handle licenses as expected. Now, Terminal Services connect as expected.

All users of rdesktop are advised to upgrade to these updated packages, which fix these bugs.

5.274. RDMA

5.274.1. RHBA-2012:1423 – rdma bug fix update

An updated rdma package that fixes a bug is 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.

Bug Fix

BZ#860943

Prior to the release of Red Hat Enterprise Linux 6.3, the kernel created the InfiniBand device files in the wrong place and a udev rules file was used to force the devices to be created in the proper place. With the update to 6.3, the kernel was fixed to create the InfiniBand device files in the proper place, and so the udev rules file was removed as no longer being necessary. However, a bug in the kernel device creation meant that, although the devices were now being created in the right place, they had incorrect permissions. Consequently, when users attempted to run an RDMA application as a non-root user, the application failed to get the necessary permissions to use the RDMA device and the application terminated. This update puts a new udev rules file in place. It no longer attempts to create the InfiniBand devices since they already exist, but it does correct the device permissions on the files.

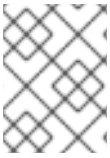
Users of rdma are advised to upgrade to this updated package, which fixes this bug.

5.275. RDMA

5.275.1. RHBA-2012:0770 – RDMA stack bug fix and enhancement update

Updated RDMA packages that fix various 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 RDMA packages have been upgraded to the latest upstream versions which provide a number of bug fixes and enhancements over the previous versions (BZ#739138).

BZ#814845

The **rdma_bw** and **rdma_lat** utilities provided by the perftest package are now deprecated and will be removed from the perftest package in a future update. Users should use the following utilities instead: **ib_write_bw**, **ib_write_lat**, **ib_read_bw**, and **ib_read_lat**.

Bug Fixes

BZ#696019

Previously, the **rping** utility did not properly join threads on shutdown. Consequently, on iWARP connections in particular, a race condition was triggered that resulted in the rping utility terminating unexpectedly with a segmentation fault. This update modifies rping to properly handle thread teardown. As a result, rping no longer crashes on iWARP connections.

BZ#700289

Previously, the kernel RDMA Connection Manager (rdmacm) did not have an option to reuse a socket port before the timeout had expired on that port after the last close. Consequently, when trying to open and close large numbers of sockets rapidly, it was possible to run out of suitable sockets that were not waiting in the timewait state. This update improves the kernel rdmacm provider to implement the SO_REUSEADDR option available for TCP/UDP sockets, which allows a socket that is closed but still in the timewait state to be reused when needed. As a result, it is now much more difficult to run out of sockets because the rdmacm provider does not need to wait for them to expire from the timewait state before they can be reused.

BZ#735954

The framework of the MVAPICH2 process manager **mpirun_rsh** in the mvapich2 package was broken. Consequently, all attempts to use mpirun_rsh failed. This update upgrades mpirun_rsh to later MVAPICH2 upstream sources that resolve the problem. As a result, mpirun_rsh works as expected.

BZ#747406

Previously, the permissions on the **/dev/ipath*** files were not permissive enough for normal users to access. Consequently, when a normal user attempted to run a Message Passing Interface (MPI)

application using the Performance Scaled Messaging (PSM) Byte Transfer Layer (BTL), it failed due to the inability to open files starting with **/dev/ipath**. This update makes sure that the files starting with **/dev/ipath** have the correct permissions to be opened in read-write mode by normal users. As a result, attempts to run an MPI application using the PSM BTL succeed.

BZ#750609

Previously, mappings from InfiniBand bit values to link speeds only extended to Quad Data Rate (QDR). Consequently, attempts to use newer InfiniBand cards that supported speeds faster than QDR did not work because the stack did not understand the bit values in the link speed field. This update adds FDR (Fourteen Data Rate), FDR10, and EDR (Enhanced Data Rate) link speeds to the kernel and user space libraries. Users can now make use of newer InfiniBand cards at these higher speeds.

BZ#754196

OpenSM did not support the **subnet_prefix** option on the command line. Consequently, in order to have two instances of OpenSM running on two different fabrics at the same time and on the same machine, the sysadmin had to edit two different **opensm.conf** files and specify the subnet_prefix separately in each file in order to have different prefixes on the different subnets. With this update, OpenSM accepts a subnet_prefix option and the OpenSM init script now starts OpenSM using this option when it is being started on multiple fabrics. As a result, a sysadmin is no longer required to hand edit multiple **opensm.conf** files to create otherwise identical configurations that only vary by which fabric they are managing.

BZ#755459

Previously, **ibv_devinfo** (a program included in libibverbs-utils) did not catch bad port numbers on the command line and return an error code. Consequently, scripts could not reliably tell whether or not the command had succeeded or failed due to a bad port number. This update fixes ibv_devinfo so that it returns a non-zero error condition when a user attempts to run it on a non-existent InfiniBand device port. As a result, scripts can now tell for certain if the port value they pass to ibv_devinfo was a valid port or was out of range.

BZ#758498

Initialization of RDMA over Converged Ethernet (RoCE) based queue pairs (QPs) was not completed successfully when initialization was done through libibverbs and not through librdmacm. Consequently, attempting to open the connection failed and the following error message was displayed:

```
cannot transition QP to RTR state
```

This updated kernel stack provides a fix for the libibverbs based RoCE QP creation and now users can properly create QPs whether they use libibverbs or librdmacm as the connection initiation method.

BZ#768109

Previously, the openmpi library did not honor the tcp_port_range settings. Consequently, if users wished to limit the **TCP** ports that openmpi used they could not do so. This update to a later upstream version that does not have this problem allows users to now limit which TCP ports openmpi attempts to use.

BZ#768457

Previously, the shared OpenType font library “libotf.so.0” was provided by both the openmpi package and the libotf package. Consequently, when an RPM spec file requested libotf.so.0 in order

to operate properly, Yum could install either openmpi or libotf to satisfy the dependency, but as these two packages do not provide compatible libotf.so.0 libraries, the program might or might not work depending on whether or not the right provider was selected. The libotf.so.0 in openmpi is not intended for other applications to link against, it is an internal library. With this update, libotf.so.0 in openmpi is excluded from RPM's library identification searches. As a result, applications linking against libotf will get the right libotf, and openmpi will not accidentally be installed to satisfy the need for libotf.

BZ#773713

There was a race condition in handling of completion events in the **perftest** programs. Under certain conditions, the perftest program being used would terminate unexpectedly with a segmentation fault. This update adds separate send receive completion queues in place of the single completion queue for both send and receive operations. The race between the finish of a send and the finish of a receive is thereby avoided. As a result, the perftest applications no longer crash with a segmentation fault.

BZ#804002

The **rds-ping** tool did not check to make sure that a socket was available before sending the next ping packet. Consequently, when the timeout between packets was set very small by the user, packets could fill up all available sockets and then overwrite one of the sockets before any ping-packets were returned. This resulted in corruption in the rds-ping data structures and eventually rds-ping terminated unexpectedly with a segmentation fault. With this update, the rds-ping program stalls on sending any more packets if there are no sockets without outstanding packets. As a result, rds-ping no longer crashes with a segmentation fault when the timeout between packets is very small.

BZ#805129

Due to a bug in the **libmlx4.conf** modprobe configuration, usage of modprobe could result in an infinite loop of modprobe processes. If the bug was encountered, the processes would continually fork until there were no processes able to run and the system would become unresponsive. This update improves the code and as a result an incorrect configuration of options in **/etc/modprobe.d/libmlx4.conf** no longer results in a system that is unresponsive and that requires a hard reboot in order to be restored to proper operation.

BZ#808673

The **qperf** application had an outdated constant for ***PF_RDS*** in its source code that did not match the officially assigned value for PF_RDS and so qperf would compile with the wrong PF_RDS constant. Consequently, when it was run it would mistakenly think RDS (Reliable Datagram Service) was not supported on the machine even when it was and would refuse to run any RDS tests. This update removes the PF_RDS constant from the qperf source code so that it will pick up the correct constant from the system header files. As a result, qperf now properly runs RDS performance tests.

BZ#815215

The **srptools** RPM did not automatically add the SCSI Remote Protocol daemon (srpd) to the service list. Consequently, the **chkconfig --list** command would not show the srpd service at all and the service could not be enabled. The srptools RPM now properly adds the srpd init script to the list of available services (it is disabled by default). Users can now see the srpd service using chkconfig --list and can enable the srpd service with the **chkconfig --level 345 srpd on** command.

BZ#815622

There was a bad test in the **rdma init** script. Consequently, the rds module would be loaded even if the user had configured it not to load. This update corrects the test in the init script so that all

conditions must be met instead of just the first condition. As a result, the rds module is only loaded when the user has configured it to be loaded or if autoloaded by the kernel due to rds usage on the local machine.

Enhancements

BZ#700285

On large InfiniBand networks, Subnet Administration service lookups consumed a large amount of bandwidth. Consequently, it could take upwards of 1 minute to look up a route from one machine to another if the network InfiniBand Subnet Manager (OpenSM) was heavily congested. This update adds the InfiniBand Communication Management Assistant (ibacm) that caches routes in a similar manner to the ARP cache for Ethernet. The ibacm program caches PathRecords from the Subnet Administration service (SA) which includes information such as MTU (Maximum Transmission Unit), SL (Service Level), SLID (Source Local Identifier) and DLID (Destination Local Identifier) for InfiniBand paths. This information is important to set up QP's properly. As a result, large subnets with many nodes will have reduced overall SA Query traffic and route lookup times.

Users of RDMA should upgrade to these updated packages, which provide numerous bug fixes and enhancements.

5.276. READLINE

5.276.1. RHBA-2012:0834 – readline bug fix and enhancement update

Updated readline packages that fix a bug and add an enhancement are now available for Red Hat Enterprise Linux 6.

The Readline library provides a set of functions that allow users to edit command lines.

Bug Fix

BZ#722942

Previously, the Readline library defined isxdigit as a macro but it is a template function in C++. Consequently, when compiling C++ code, every other C++ header had to be included before the Readline headers in order to avoid compiling errors. With this update, isxdigit is no longer defined as a macro if the compiled program is in C++. As a result, developers no longer have to avoid including Readline headers first when working with C++.

Enhancement

BZ#244350

This update adds advisory TTY input audit events, to record the lines actually reported to applications.

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

5.277. REDHAT-RELEASE

5.277.1. RHEA-2012:0971 – redhat-release enhancement update for Red Hat Enterprise Linux 6.3

An enhanced redhat-release package is now available for Red Hat Enterprise Linux 6.3.

The redhat-release package contains licensing information regarding, and identifies the installed version of, Red Hat Enterprise Linux.

This updated redhat-release package reflects changes made for the release of Red Hat Enterprise Linux 6.3.

Users of Red Hat Enterprise Linux 6 are advised to upgrade to this updated redhat-release package, which adds this enhancement.

5.278. REDHAT-RPM-CONFIG

5.278.1. RHBA-2012:0911 – redhat-rpm-config bug fix and enhancement update

An updated redhat-rpm-config package that fixes several bugs and adds an enhancement is now available for Red Hat Enterprise Linux 6.

The redhat-rpm-config package is used during building of RPM packages to apply various default distribution options determined by Red Hat. It also provides a few Red Hat RPM macro customizations, such as those used during the building of Driver Update packages.

Bug Fixes

BZ#680029

Previously, the %kernel_module_package macro did not handle the "-v" and "-r" optional version and release override parameters correctly. Consequently, the specified version and release number were not used when the RPM package was built. This bug has been fixed and these parameters are now handled properly.

BZ#713638

Previously, a script, which generates "modalias"-style dependencies for Driver Update packages in Red Hat Enterprise Linux 6, was not executable and thus could not function properly. This bug has been fixed and these dependencies are now generated as expected.

BZ#713992

When the kabi-whitelists package is installed, the %kernel_module_package macro did not automatically perform a check against the Red Hat kernel ABI interface (kABI). Consequently, when a package was being built, the macro did not warn when the resulting modules used kernel symbols that were exported but not part of the kABI. With this update, the abi_check.py script has been added to perform the check and return a warning during the build process if kabi-whitelists is not installed, thus fixing this bug.

BZ#767738

In certain cases, the dependency-generation scripts that produce information about automatic kernel symbol during the build process of a Driver Update package generated incorrect dependencies. This bug has been fixed and dependencies are now generated correctly.

Enhancement

BZ#652084

The path of the autoconf configuration script invoked by the %configure macro can now be customized by overriding the %_configure macro. In addition, this can be of use when building out-of-tree packages.

Users of redhat-rpm-config are advised to upgrade to this updated package, which fixes these bugs and adds this enhancement.

5.279. RED HAT ENTERPRISE LINUX RELEASE NOTES

5.279.1. RHEA-2012:0979 – Red Hat Enterprise Linux 6.3 Release Notes

Updated packages containing the Release Notes for Red Hat Enterprise Linux 6.3 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 6.3 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 6.3 Release Notes:

https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/6.3_Release_Notes/index.html

5.280. RESOURCE-AGENTS

5.280.1. RHBA-2012:1419 – resource-agents bug fix update

Updated resource-agents packages that fix a bug 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 the Pacemaker and rgmanager service managers.

Bug Fix

BZ#864364

If the contents of the /proc/mounts file changed during a status check operation of the fs.sh file system resource agent, the status check could incorrectly detect that a mount was missing and mark a service as failed. This bug has been fixed and fs.sh no longer reports false failures in the described scenario.

All users of resource-agents are advised to upgrade to these updated packages, which fix this bug.

5.280.2. RHBA-2012:1515 – resource-agents bug fix update

Updated resource-agents packages that fix a bug 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 (HA) environment for both the Pacemaker and rgmanager service managers.

Bug Fix

BZ#878023

Previously, when device failures caused logical volumes to go missing, HA LVM was unable to shut down. With this update, services can migrate to other machines that still have access to the devices, thus preventing this bug.

All users of resource-agents are advised to upgrade to these updated packages, which fix this bug.

5.280.3. RHBA-2012:0947 – resource-agents bug fix and enhancement update

Updated resource-agents packages that fix multiple bugs and add three 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#728086

Prior to this update, the fs-lib.sh resource agent library ignored the error codes greater than '1'. As a consequence, fs-lib.sh failed to recognize errors when a mount returned an error with a different error code, for example an iSCSI mount. This update modifies the underlying code so that the fs-lib.sh resource agent library now recognizes all errors as expected.

BZ#742859

Prior to this update, the Apache resource agent did not correctly generate the IPv6 configuration for the configuration file. As a consequence, Apache failed to work with IPv6 addresses. This update modifies the underlying code so that the Apache resource agent now generates a valid configuration file when IPv6 is in use.

BZ#746996

Prior to this update, the SAP Web Dispatcher and the TREX service were not monitored in the SAP resource agent script. This update adds the SAP Web Dispatcher and the TREX Service to the list of services that are checked for SAP. Now, the SAP Web Dispatcher and the TREX Service are monitored.

BZ#749713

Prior to this update, missing etab entries were not recreated due to an error in a regular expression and an incorrect flag on the "clufindhostname" command. As a consequence, NFS exports were not automatically recovered. This update corrects the regular expression and uses the "clufindhostname" command as expected. Now, NFS exports recover automatically when entries are removed from the etab file.

BZ#784357

Prior to this update, the configuration path variable for the resource agent was not correctly set. As a consequence, the wrong path for configuration files was used. This update modifies the configuration path variable so that the common configuration directory is now correctly set to prevent problems with the resource agents for Samba, Apache, and others.

BZ#799998

Prior to this update, the "netfs" script did not identify whether the file systems to be checked were network file systems before denying multiple mounts. As a consequence, network file systems could not be added twice. This update modifies the "netfs" script so that it verifies if file systems are network file systems and allows multiple mounts for these. Now, multiple mounts of the same network file system are allowed.

Enhancements

BZ#712174

Prior to this update, no option to set tunnelled migrations with the Kernel-based Virtual Machine (KVM) was available. This update adds the "--tunnelled" option to the vm.sh resource agent to allow encrypted migrations between qemu virtual machines.

BZ#726500

Prior to this update, the SAP resource agent scripts did not reflect changes in the upstream version. This update merges Pacemaker and the Heartbeat SAP resource agent with the upstream version.

BZ#784209

The SAP database resource agent has been synchronized with the upstream resource agent to provide additional functionality and bug fixes.

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

5.281. RGMANAGER

5.281.1. RHBA-2012:0897 – rgmanager bug fix and enhancement update

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

The rgmanager packages contain the Red Hat Resource Group Manager, which provides the ability to create and manage high-availability server applications in the event of system downtime.

Bug Fixes

BZ#635152

Previously, rgmanager incorrectly called the rg_wait_threads() function during cluster reconfiguration. This could lead to an internal deadlock in rgmanager which caused the cluster services to become unresponsive. This irrelevant call has been removed from the code and deadlocks now no longer occur during cluster reconfiguration.

BZ#727326

When enabling a service using the clusvcadm command with the "-F" option, rgmanager did not update the service owner information before responding to clusvcadmn. Consequently, clusvcadm could print incorrect information about which cluster node the service was running on. This update modifies rgmanager to update the owner information prior to responding to clusvcadm, and the command now provides the correct information.

BZ#743218

Under certain circumstances, a "stopped" event could be processed after a service and its dependent services had already been restarted. This forced the dependent services to restart erroneously. This update allows rgmanager to ignore the "stopped" events if dependent services have already been started, and the services are no longer restarted unnecessarily.

BZ#744824

Resource Group Manager did not handle certain inter-service dependencies correctly. Therefore, if a service was dependent on another service that was running on the same cluster node, the dependent service became unresponsive during the service failover and remained in the recovering state. With this update, rgmanager has been modified to check a service state during failover and stop the service if it is dependent on the service that is failing over. Resource Group Manager then tries to start this dependent service on other nodes as expected.

BZ#745226

The "-F" option of the clusvcadm command allows rgmanager to start a service according to failover domain rules. This option was not previously described in the command's manual pages. With this update, the "-F" option has been properly documented in the clusvcadm(8) manual page.

BZ#796272

Previously, if a newly added service failed to start on the first cluster node, rgmanager could try to relocate the service to another cluster node before the cluster configuration was updated on that node. Consequently, the service was set to the "recovering" state and had to be manually re-enabled in order to start. This update modifies rgmanager to retry the relocation process until after the cluster configuration has been updated on the node. The service can now be relocated as expected.

BZ#803474

Due to an invalid pointer dereference, rgmanager could terminate unexpectedly with a segmentation fault when central processing mode was enabled on a cluster node. With this update, the pointer dereference has been corrected, and rgmanager no longer crashes when central processing mode is enabled.

BZ#807165

Previously, in central processing mode, rgmanager failed to restart services that depended on a service that failed and was recovered. With this update, during the recovery of a failed service, any services that depend on it are restarted.

Enhancement

BZ#799505

This update introduces a feature which enables rgmanager to utilize Corosync's Closed Process Group (CPG) API for inter-node locking. This feature is automatically enabled when Corosync's Redundant Ring Protocol (RRP) feature is enabled. Corosync's RRP feature is considered fully supported. However, when used with the rest of the High-Availability Add-Ons, it is considered a Technology Preview.

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

5.282. RHN-CLIENT-TOOLS AND YUM-RHN-PLUGIN

5.282.1. RHBA-2012:0752 – rhn-client-tools and yum-rhn-plugin bug fix and enhancement update

Updated rhn-client-tools and yum-rhn-plugin packages that fix several bugs and add two enhancements are now available for Red Hat Enterprise Linux 6.

The rhn-client-tools and yum-rhn-plugin packages provide programs and libraries that allow a system to receive software updates from Red Hat Network or Red Hat Network Satellite.

Bug Fixes

BZ#729342

When a client was set up to connect to a server using the *Secure Sockets Layer* (SSL) protocol and the validation of the server's SSL certificate failed, an **SSLCertificateVerifyFailedError** exception was raised and a traceback was written to the **/var/log/up2date** file. This update corrects the exception handling mechanism to ensure that a proper error message is displayed in this scenario.

BZ#735339

The previous version of yum-rhn-plugin did not properly clean its cache, which may have caused various misleading messages to be displayed during subsequent Yum operations. To prevent this behavior, this update ensures that yum-rhn-plugin cleans its cache properly.

BZ#744709

When using the **firstboot** application to register the system with RHN Classic, entering incorrect credentials previously resulted in a traceback. With this update, an explanatory error message is displayed in this situation.

BZ#745095

When a virtual host was registered with Red Hat Network Satellite and the user attempted to register a virtual guest, the **rhn_register** utility did not associate the guest with the virtual host. Consequent to this, the virtual host appeared to have no virtual guests. This update ensures that **rhn_register** correctly pairs the guest with the managed virtual host.

BZ#746983

When a client was banned due to abuse of service, an attempt to run the previous version of the **rhn-channel** utility failed with both a traceback and an error message. With this update, the utility now displays only the error message.

BZ#748876

Previously, running the **rhnreg_ks** utility with a relative path to the SSL certificate caused it to store this relative path in the configuration file. Consequently, an attempt to run any other utility that uses this configuration file from a different directory rendered such a utility unable to open the certificate. This update adapts **rhnreg_ks** to store an absolute path.

BZ#751893

The web user interface of Red Hat Network Satellite now displays the BIOS version of managed Red Hat Enterprise Linux 6 clients, as expected.

BZ#768045

When **Red Hat Subscription Manager** was installed but the user decided to register the system with RHN Classic, Red Hat Subscription Manager incorrectly reported that the system in non-compliant.

With this update, when the user registers a system with RHN Classic, the **rhnreg_ks** and **rhn_register** tools now notify Red Hat Subscription Manager over D-Bus.

BZ#771167

When the user attempted to register a system with RHN Classic and provided invalid credentials, the **rhn_register** utility reported an error at the very end of the registration process. This update adapts **rhn_register** to report invalid credentials immediately.

BZ#781421

Previously, the presence of Unicode characters in a localized error message caused an additional traceback to be displayed. With this update, error messages are encoded in UTF-8 before they are printed to standard error.

BZ#788903

Prior to this update, when the user disabled a channel in the **/etc/yum/pluginconf.d/rhnplugin.conf** configuration file, a subsequent update of the yum-rhn-plugin package re-enabled this channel. This update ensures that such configuration changes are persistent.

BZ#799926

Due to a new layout of the Python modules for **Red Hat Subscription Manager**, the **firstboot** application did not skip the login screen when the system was already registered. This update corrects this error.

BZ#809241

Due to an error in the list of required packages, an attempt to run the **rhn-setup** utility on a system without the newt-python package installed could fail with an **ImportError**. This update adds newt-python as a dependency of the rhn-client-tools package.

BZ#817567

Under certain circumstances, the **rhn_check** utility could incorrectly report an attempt to update already updated packages as failed. To prevent this, yum-rhn-plugin has been adapted to ensure that packages that are already updated are properly removed from the list of packages that are scheduled for an update.

BZ#735346

The **rhn_check** utility no longer displays debugging messages without a space after the **do_call** keyword.

BZ#751292, BZ#759786

When the network connection is reset by peer or a general network error occurs, yum-rhn-plugin now displays an explanatory error message and no longer fails with a traceback.

Enhancements

BZ#569790, BZ#749281, BZ#767679

The rhn-client-tools packages now support IPv6.

BZ#772070

The user interface of the firstboot application has been redesigned to improve the usability of the system registration screen.

All users of rhn-client-tools and yum-rhn-plugin are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.282.2. RHBA-2013:1385 – yum-rhn-plugin bug fix update

Updated yum-rhn-plugin packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

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#993105

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.

5.283. RICCI

5.283.1. RHBA-2012:0898 – ricci bug fix and enhancement update

Updated ricci packages that fix multiple 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#724014

Previously, the ccs utility was not able to configure IPV6 hosts. The ccs utility has been updated to support the IPV6 protocol, so that it can be used to configure IPV6 hosts.

BZ#726772

The cman and ccs utilities contain different cluster schemas for Red Hat Enterprise Linux 6.2. A different cluster schema on a node from the schema contained in ccs could cause a valid cluster.conf file to be interpreted as invalid by ccs. A new ccs option, "--getschema", has been added. Using this option, if cluster schemas are different between what is included in ccs and the information on a node, the cluster schema on the node is used.

BZ#731113

Previously, the ccs utility did not check to see if the cluster.conf file was invalid. Therefore, when parsing an invalid cluster.conf file, ccs could terminate unexpectedly followed by a traceback, however with no detailed information about the problem provided. With this update, ccs checks to

see whether cluster.conf is valid; if the file is invalid, ccs prints a helpful error message and exits gracefully.

BZ#738008

Previously, the "ccs_sync" command did not return a non-zero exit code if an error occurred or the ricci daemon was not running, even when running the command with the "-w" option (to exit with a failure status if any warnings were issued). The underlying source code has been modified so that "ccs_sync" with the "-w" option now returns "1" on failure.

BZ#738567

When running the "ccs" command with the "--checkconf" and "-f" options, ccs verifies that all the nodes in the file specified contain the same cluster.conf file. The XML code in the configuration file is indented whereas the XML code in the live configuration file from ricci is not, and therefore bare comparison of such strings always failed. As a consequence, ccs did not verify each node in the local cluster.conf file. The comparison method has been improved so that "ccs --checkconf -f" now correctly verifies whether all the nodes in the file contain the same cluster.conf file.

BZ#742345

Previously, if the user installed the ricci packages but did not install the modcluster package, any attempt to run cluster commands using ricci failed. With this update, users are no longer allowed to install only the ricci packages; the modcluster package is now required as a dependency. As a result, cluster commands can be executed as expected.

BZ#770637

Previously, when the user configured a virtual machine service, the virtual machine was not displayed in the output of the "ccs --lsservices" command. The ccs utility has been modified to specifically check for virtual machine services in the configuration file. Now, running "ccs" with the "--lsservices" option prints the proper output.

BZ#773383

Previously, if the user installed ricci and did not set a user password, the user was unable to connect to ricci. This could lead to confusion, because the user was not prompted to provide the password and therefore not aware of the requirement. Also, no explanation was logged in the system log. With this update, ricci logs a warning message to syslog on startup if the ricci user password is not set.

Enhancements

BZ#738797

Prior to this update, the ricci daemon processed tasks, but did not log any information about when the tasks were run. With this update, when ricci is asked to spawn a worker process, the date and time information is now logged using syslog.

BZ#758823

Prior to this update, the ccs utility did not provide a way to configure Redundant Ring Protocol (RRP). This update adds additional configuration options for RRP into the ccs utility so that users can configure RRP using ccs.

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

5.284. RPCBIND

5.284.1. RHEA-2012:0974 – rpcbind enhancement update

Enhanced rpcbind packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

The rpcbind utility maps RPC (Remote Procedure Call) services to the ports on which the services listen and allows the host to make RPC calls to the RPC server.

Enhancement

BZ#726954

The rpcbind tool did not drop supplementary groups and the groups remained available after rpcbind downgraded from root-group privileges. As a security hardening measure, rpcbind now drops root privileges correctly, running as a non-root user after it has bound to its privileged network port.

Users of rpcbind are advised to upgrade to these enhanced packages, which add this enhancement.

5.284.2. RHBA-2013:1453 – rpcbind bug fix

Updated rpcbind packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The rpcbind utility maps RPC (Remote Procedure Call) services to the ports on which the services listen and allows the host to make RPC calls to the RPC server.

Bug Fix

BZ#858573

Previously, in the insecure mode, which enables non-root users to set or unset ports, a privileged port was required. As only root users can obtain a privileged port, non-root users could not set or unset ports. To fix this bug, the privileged port has been removed, and thus non-root users are now allowed to set or unset ports on the loopback interface.

All users of rpcbind are advised to upgrade to these updated packages, which fix this bug.

5.285. RPMDEVTOOLS

5.285.1. RHBA-2012:1313 – rpmdevtools bug fix update

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

The rpmdevtools packages contain scripts and (X)Emacs support files to aid in development of RPM packages.

Bug Fix

BZ#730770

Prior to this update, the sample spec files referred to a deprecated BuildRoot tag. The tag was ignored if it was defined. This update removes the BuildRoot tags from all sample spec files.

All users of rpmdevtools are advised to upgrade to these updated packages, which fix this bug.

5.286. RPM

5.286.1. RHBA-2012:0909 – rpm bug fix and enhancement update

Updated rpm packages that fix multiple bugs and add various enhancements 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#799317

Previously, presence of ELF files prevented cross-architecture obsoletion of packages on multi-arch systems, causing file conflicts when a supposedly obsoleted package was not removed. With this update, obsoletes are now processed for all matching package names regardless of their contents, allowing scenarios like eliminating no longer needed 32-bit package variant on 64-bit multi-arch systems to work as expected.

BZ#746190

Previously, a bug in execution of package scriptlets utilizing RPM's embedded Lua interpreter could have caused RPM's working directory to change inadvertently, resulting in a failure to install remaining local packages in the transaction unless absolute paths were used to address the packages on the "rpm" or "yum" command line. With this update, the Lua scriptlet execution now always saves and restores the current working directory, ensuring correct operation regardless of whether absolute or relative paths to packages are used.

BZ#785236

Previously, the "-D" shortcut option for "--define" was incorrectly taken as a shortcut for "--predefine", which led to incorrect macro evaluation when attempting to override macros from system configuration. The "-D" shortcut option now equals "--define" as intended and documented.

BZ#768516

Previously, RPM's "--last" query format output could have been ambiguous on multi-arch systems such as AMD64/Intel 64 as package architecture was omitted. Package's architecture is now included in "--last" output as well, making it non-ambiguous and also consistent with the default query output format.

BZ#664427

As the build dependencies recorded in source packages can vary depending on the architecture where the source packages happened to be generated, using the yum-builddep utility on a source package does not always report correct results. RPM's Python bindings have now been enhanced to permit yum-builddep to operate on spec files directly, ensuring that the correct build dependency information for the local system is used.

BZ#752119

Previously, certain multi-line brace constructs could have caused the automatic Perl dependency generator script to miss dependencies from pe. The generator has now been updated to properly handle these situations.

Enhancements

BZ#714678

When building packages on file systems with a very high number of files, the on-disk inode numbers could have been truncated in RPM's 32bit-integer-based hardlink tracking, resulting in incorrect package generation and, consequently, installation. RPM now uses per-package virtual numbering for hardlink tracking to eliminate the possibility of truncation, ensuring correct operation regardless of physical inode numbers at package build time, in a backwards compatible way.

BZ#736960

Previously, RPM ignored any exit codes from %pretrans package scriptlets. This was inconsistent with semantics of other scriptlets and prevented the possibility of early abort of package installation before the transaction really starts. RPM now treats %pretrans failure similarly to that of %pre: the package with failing %pretrans scriptlet is not installed at all.

BZ#761000

Packages for Fedora 17 or later require a special rpmlib() dependency provided by RPM to track the /usr merge that was completed in Fedora 17, otherwise it will no be possible to use, for example, mock chroot to install and build packages for that distribution. This special tracking dependency has been added to RPM now to allow Red Hat Enterprise Linux 6 to be used as a host for building packages for these newer Fedora versions.

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

5.287. RSYNC

5.287.1. RHBA-2012:0473 – rsync bug fix update

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

The rsync tool is used to copy and synchronize files locally or across a network. The rsync works very fast because it uses delta encoding and sends only differences in files instead of whole files. The rsync is also used as a powerful mirroring tool.

Bug Fixes

BZ#737539

When creating a sparse file that was zero blocks long, the "rsync --sparse" command did not properly truncate the sparse file at the end of the copy transaction. As a consequence, the file size was bigger than expected. With this update, the underlying source code has been modified to ensure proper truncating of such files.

BZ#804916

Previously, the rsync utility could terminate unexpectedly with the following error during a data transfer:

```
Inflate (token) returned -5
```

This happened if the block size was exactly of the size of the CHUNK_SIZE constant. The output buffer was completely filled after calling the inflate() function for the first time, and it was therefore

not possible to obtain remaining buffer output when calling inflate() the next time. The Z_BUF_ERROR constant is now handled properly, and so prevents rsync from terminating in the described scenario.

All users of rsync are advised to upgrade to this updated package, which fixes these bugs.

5.287.2. RHBA-2013:1501 – rsync bug fix update

Updated rsync packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The rsync tool is used to copy and synchronize files locally or across a network. The rsync works very fast because it uses delta encoding and sends just differences in files instead of whole files. The rsync is also used as powerful mirroring tool.

Bug Fix

BZ#1022357

Previously, the rsync tool did not check whether the inbuf variable is non-empty. As a consequence, rsync terminated unexpectedly while trying to do the required encoding in a loop. With this update, rsync checks whether inbuf is non-empty and no longer crashes in the described scenario.

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

5.288. RSYSLOG

5.288.1. RHSA-2012:0796 – Moderate: rsyslog security, bug fix, and enhancement update

Updated rsyslog packages that fix one security issue, multiple 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. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

The rsyslog packages provide an enhanced, multi-threaded syslog daemon.

Security Fix

CVE-2011-4623

A numeric truncation error, leading to a heap-based buffer overflow, was found in the way the rsyslog imfile module processed text files containing long lines. An attacker could use this flaw to crash the rsyslogd daemon or, possibly, execute arbitrary code with the privileges of rsyslogd, if they are able to cause a long line to be written to a log file that rsyslogd monitors with imfile. The imfile module is not enabled by default.

Bug Fixes

BZ#727380

Several variables were incorrectly deinitialized with Transport Layer Security (TLS) transport and keys in PKCS#8 format. The rsyslogd daemon aborted with a segmentation fault when keys in this format were provided. Now, the variables are correctly deinitialized.

BZ#756664

Previously, the imgssapi plug-in initialization was incomplete. As a result, the rsyslogd daemon aborted when configured to provide a GSSAPI listener. Now, the plug-in is correctly initialized.

BZ#767527

The fully qualified domain name (FQDN) for the localhost used in messages was the first alias found. This did not always produce the expected result on multihomed hosts. With this update, the algorithm uses the alias that corresponds to the hostname.

BZ#803550

The gtls module leaked a file descriptor every time it was loaded due to an error in the GnuTLS library. No new files or network connections could be opened when the limit for the file descriptor count was reached. This update modifies the gtls module so that it is not unloaded during the process lifetime.

BZ#805424

rsyslog could not override the hostname to set an alternative hostname for locally generated messages. Now, the local hostname can be overridden.

BZ#807608

The rsyslogd init script did not pass the lock file path to the 'status' action. As a result, the lock file was ignored and a wrong exit code was returned. This update modifies the init script to pass the lock file to the 'status' action. Now, the correct exit code is returned.

BZ#813079

Data could be incorrectly deinitialized when rsyslogd was supplied with malformed spool files. The rsyslogd daemon could be aborted with a segmentation fault. This update modifies the underlying code to correctly deinitialize the data.

BZ#813084

Previously, deinitialization of non-existent data could, in certain error cases, occur. As a result, rsyslogd could abort with a segmentation fault when rsyslog was configured to use a disk assisted queue without specifying a spool file. With this update, the error cases are handled gracefully.

BZ#820311

The manual page wrongly stated that the '-d' option to turn on debugging caused the daemon to run in the foreground, which was misleading as the current behavior is to run in the background. Now, the manual page reflects the correct behavior.

BZ#820996

rsyslog attempted to write debugging messages to standard output even when run in the background. This resulted in the debugging information being written to some other output. This was corrected and the debug messages are no longer written to standard output when run in the background.

BZ#822118

The string buffer to hold the distinguished name (DN) of a certificate was too small. DNs with more than 128 characters were not displayed. This update enlarges the buffer to process longer DNs.

Enhancements

BZ#672182

Support for rate limiting and multi-line message capability. Now, rsyslogd can limit the number of messages it accepts through a UNIX socket.

BZ#740420

The addition of the "/etc/rsyslog.d/" configuration directory to supply syslog configuration files.

All users of rsyslog are advised to upgrade to these updated packages, which upgrade rsyslog to version 5.8.10 and correct these issues and add these enhancements. After installing this update, the rsyslog daemon will be restarted automatically.

5.289. RUSERS

5.289.1. RHBA-2012:0404 – rusers bug fix update

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

The rusers program allows users to find out who is logged into various machines on the local network. The rusers command produces output similar to output of the who utility, but for the specified list of hosts or for all machines on the local network.

Bug Fix

BZ#697862

Previously, no dependency on the rpcbind package was specified in the rstatd and rusers SysV init scripts. In addition, when the rstatd and rusersd services were started, a check to see if networking was enabled was not performed, and an incorrect exit code was returned. This update adds rpcbind dependency to the rusersd and rstatd init scripts. Also, the SysV init scripts have been adjusted to return correct exit codes. Checks are now performed when starting the rstatd and rusersd services to see whether networking is available and binding to rpcbind was successful.

All users of rusers are advised to upgrade to these updated packages, which fix this bug.

5.290. S390UTILS

5.290.1. RHBA-2012:0885 – s390utils bug fix and enhancement update

Updated s390utils packages that fix several bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The s390utils packages contain utilities related to Linux for the IBM System z architecture.

Bug Fixes

BZ#747937

Previously, when calculating memory size on systems with less than 256 MB or more than 256 GB of memory, an "Illegal division by zero" error occurred. The code has been fixed and as a result, memory device size calculation now works correctly.

BZ#769416

When the **ttyrun** tool fails to open a terminal device, an error condition is reported to syslog. Depending on the environment, multiple terminal devices might not be available and, hence, multiple error messages are issued. This update introduces a verbose option to switch on or off syslog messages issued by ttyrun.

BZ#772576

The HiperSockets Network Concentrator **xcec-bridge** utility issues a warning if the length of a send buffer is different from the specified length of the packet to be sent. The condition check was wrong. This update corrects the code and checking of the packet length to be sent now works correctly.

BZ#783161

Calling the **zipl** utility changes data in the file system but also writes to the device node. The built-in sync call only flushed file system buffers which did not ensure that data on the device node /dev/dasdX had been written. This could cause an incorrect bootmap to be used during *Initial Program Load* (IPL) and the wrong kernel to be loaded or result in an unbootable system. This update adds a call to the fsync() function for all writable opened file descriptors. As a result, all the relevant data is written to disk and the problem is solved.

BZ#785761

The **qethconf** utility searched for an exact match of an **IPv6** address. Consequently, an IPv6 address written in capital letters was not deleted. Thus update makes the IPv6 search case insensitive and the problem no longer occurs.

BZ#785762

A service fix added to z/VM 6.1 broke the interface of the *Control Program* (CP) command for querying a network interface (NIC) and the layer 3 address was not detected automatically under z/VM 6.x. Consequently, configuring a network interface via the znetconf tool failed. This update corrects the problem and the CP command for querying a virtual network interface now works as expected.

BZ#797937

When mounting a disk, the **cmsfs-fuse** utility tries to use the **mmap** function on the whole disk. This can fail for very large disks or if the virtual memory is limited. This update adds a fallback to the **pread** and **pwrite** commands in case the mmap system call fails.

BZ#800462

Parsing logic only detected the ***NONROUTER*** and ***PRIROUTER*** option for layer-3 VSWITCHes but not plain IP VSWITCH configurations. Therefore the layer 3 address for a network interface (NIC) was not detected automatically for virtual NICs connected to plain IP VSWITCHes and the virtual NIC was treated as a layer-2 device. Consequently, configuring a virtual NIC using the **znetconf** tool failed. With this update, the **IP** option is interpreted as a layer-3 indicator and configuring a virtual NIC using znetconf works as expected.

BZ#809510

The monitor record header for the stop record was incomplete. Consequently, unused monitor records of stopped processes were not stopped correctly and continued to show up in the z/VM monitor stream. This update corrects the problem and a complete header is now specified.

BZ#814311

The **zipl** helper script calls external programs without an absolute path. Consequently, when zipl is called with an empty PATH environment variable, the zipl helper script that is called by zipl is not able to execute external programs. With this update, the zipl helper script sets the PATH variable so that external programs can be located and executed as expected.

Enhancements

BZ#632347

Automatic calculation of the boot device RAM disk address for System z has been added. With this update, in place of the default address for the initial RAM disk the address is now automatically calculated using the image size. This avoids a possible file overlap for large kernel images.

BZ#633532

Enhancements to the configuration tool for IBM System z network devices have been made. With this update, the System z **qethconf** tool provides information messages when an attribute did not change as expected.

BZ#738863

The Linux kdump framework has been ported to Linux on System z and is now integrated into the existing System z stand-alone dump tools and shutdown actions framework. This leads to the following enhancements for System z kernel dumps:

- Dump time and size can be reduced using page filtering with the makedumpfile tool.
- Dump disk space sharing is possible for server farms using network dump.
- Dump setup is made easier using existing **kdump** setup GUIs of Linux distributions.

BZ#738870

Enhanced *Direct Access Storage Device* (DASD) statistics for *Parallel Access Volume* (PAV) and *High Performance FICON* (HPF) has been added. This enhancement adds the **dasdstat** tool that provides user-friendly access to the enhanced statistics provided by the kernel via the debugfs interface. It enables improved diagnosis of PAV and HPF environments to analyze and tune the DASD performance in a system, for example to give recommendations on the number of alias devices or the usage of Hyper PAV versus Base PAV. For more information about this enhancement, refer to https://www.ibm.com/developerworks/linux/linux390/documentation_dev.html.

BZ#738873

IPv6 support to the **qetharp** tool has been added. This enhancement adds IPv6 support to the qetharp tool for inspection and modification of the **ARP** cache of Open Systems Adapter (OSA) cards or HiperSockets (real and virtual) operated in layer-3 mode.

Users of s390utils should upgrade to these updated packages, which fix these bugs and add these enhancements.

5.291. SAMBA

5.291.1. RHBA-2012:0850 – samba bug fix update

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

Samba is an open-source implementation of the Server Message Block (SMB) and Common Internet File System (CIFS) protocol, which allows PC-compatible machines to share files, printers, and other information.

Bug Fixes

BZ#753143

When using Samba with the "password server" configuration setting and when the given name for that parameter was a hostname that resolved to multiple IP addresses, Samba did not correctly handle the returned addresses. Consequently, Samba failed to use one of the password servers and terminate unexpectedly. This update fixes Samba to correctly process multiple IP addresses when using a hostname with the "password server" parameter. Samba now works correctly with multiple IP addresses in the scenario described.

BZ#753747

When Samba was configured to operate in an Active Directory (AD) environment it sometimes created invalid DNS SRV queries. This happened when an empty sitename was used to compose the SRV record search string. Consequently, Samba-generated log files contained many DNS related error messages. Samba has been fixed to always generate a correct DNS SRV query and the DNS-related error message no longer occur.

BZ#755347

The smbclient tool sometimes failed to return the expected exit status code; it returned 0 instead of 1. Consequently, using smbclient in a script caused some scripts to fail. With this update, an upstream patch has been applied and smbclient now returns the correct exit status.

BZ#767656

Previously, the Winbind IDMAP interface cache did not expire as specified in the smb.conf file. Consequently, the positive and negative entries in the cache would not expire until the opposite type of query was made. This update contains a backported fix for the problem. As a result, the idmap cache time and idmap negative cache time directives now work as expected.

BZ#767659

When calling "getent passwd" for a user who had no UID, if winbind was joined to the domain with idmap_ad specified as the backend, enumerating users was enabled, and most of the users had UIDs, the enumeration stopped and the following error was displayed:

```
NT_STATUS_NONE_MAPPED
```

This update implements an upstream patch to correct the problem. As a result, if a user cannot be mapped, winbind no longer stops but continues enumerating users in the scenario described.

BZ#771812

Samba sometimes generated many debug messages such as "Could not find child XXXX -- ignoring" that were written to syslog. Consequently, although these messages are not critical, syslog could be flooded by the large amount of these messages. Samba has been fixed to no longer issue this

message to syslog automatically and syslog is no longer flooded by these samba debug messages.

BZ#788089

The pam_winbind utility used an undocumented PAM_RADIO_TYPE message which has no documented semantics. This caused the login manager gdm to terminate unexpectedly when pam_winbind was used on the system. Consequently, users could not log in when using pam_winbind. Samba has been fixed to not use the PAM_RADIO_TYPE message. Users can now use pam_winbind for authentication in GDM.

BZ#808449

Newer versions of Windows could not properly set Access Control Lists (ACLs) on a Samba share. The users were receiving an "access denied" warning. Consequently, administrators or users could not fully control ACLs on a Samba share. This update fixes the problem in Samba and ACLs can now be used as expected.

BZ#816123

An update of the system Kerberos library to a recent version made Samba binaries and libraries suddenly unusable because Samba was using a private library symbol. Consequently, Samba was no longer usable after a Kerberos update. This update corrects Samba to no longer use that private symbol. Samba now continues to operate when the Kerberos library has been updated.

All users of samba are advised to upgrade to these updated packages, which fix these bugs.

5.292. SANLOCK

5.292.1. RHEA-2012:0996 – sanlock enhancement update

Updated sanlock packages that add multiple enhancements are now available for Red Hat Enterprise Linux 6.

The sanlock packages provide a shared disk lock manager that uses disk paxos to manage leases on shared storage. Hosts connected to a common Storage Area Network (SAN) can use sanlock to synchronize the access to the shared disks. Both libvirt and vdsm can use sanlock to synchronize access to shared virtual machine (VM) images.

The sanlock packages have been upgraded to the latest upstream version, which provides a number of enhancements over the previous version. (BZ#782600)

All users of sanlock are advised to upgrade to these updated packages, which add these enhancements.

5.293. SBLIM-CIM-CLIENT2

5.293.1. RHSA-2012:0987 – Low: sblim-cim-client2 security update

Updated sblim-cim-client2 packages that fix one security issue 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 for each vulnerability from the CVE link(s) associated with each description below.

The SBLIM (Standards-Based Linux Instrumentation for Manageability) CIM (Common Information Model) Client is a class library for Java applications that provides access to CIM servers using the CIM Operations over HTTP protocol defined by the DMTF (Distributed Management Task Force) standards.

Security Fix

CVE-2012-2328

It was found that the Java HashMap implementation was susceptible to predictable hash collisions. SBLIM uses HashMap when parsing XML inputs. A specially-crafted CIM-XML message from a WBEM (Web-Based Enterprise Management) server could cause a SBLIM client to use an excessive amount of CPU. Randomization has been added to help avoid collisions.

All users of sblim-cim-client2 are advised to upgrade to these updated packages, which contain a backported patch to resolve this issue.

5.294. SCSI-TARGET-UTILS

5.294.1. RHEA-2012:0970 – scsi-target-utils enhancement update

An updated scsi-target-utils package that adds two enhancements is now available for Red Hat Enterprise Linux 6.

The scsi-target-utils package contains 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.

Enhancements

BZ#747510

Previously, scsi-target-utils did not accurately report certain parameters, such as the native block size of the "direct-store" backing store, lowest aligned logical block address, and optimal transfer length. This could decrease performance of iSCSI devices. This update adds support for the "B0 VPD" and "READ CAPACITY(16)" SCSI commands to match the characteristics of the iSCSI device to its underlying backing store, which improves performance of the iSCSI device.

BZ#605925

With this update, scsi-target-utils now supports authentication for target discovery. It may be configured using the "incomingdiscoveryuser" and "outgoingdiscoveryuser" directives in the "/etc/tgt/targets.conf" file. Configuration details can be seen in the targets.conf(5) manual page.

All users of scsi-target-utils are advised to upgrade to this updated package, which adds these enhancements.

5.295. SDL

5.295.1. RHBA-2012:0446 – SDL bug fix update

An updated SDL package that fixes multiple bugs is now available for Red Hat Enterprise Linux 6.

Simple DirectMedia Layer (SDL) is a cross-platform multimedia library designed to provide fast access to the graphics frame buffer and audio device.

Bug Fixes

BZ#733605

If the SDL_VM_GrabInput() function was called when the window of an SDL application was not visible (for example, the window was displayed on a different workspace or outside of the screen borders), the SDL library, and therefore the current application thread, was unresponsive until the window became visible and the input could be grabbed. The SDL_VM_GrabInput() function has been adjusted to return immediately with the proper error code signaling that the grab failed.

BZ#678569

Previously, calling the SDL_BlitSurface() function on overlapping rectangles when running the Streaming SIMD Extensions 3 (SSE3) optimized standard C library caused the bitmap content on the screen to be corrupted. The internal SDL_BlitCopyOverlap() function has been fixed to copy bitmaps between overlapping areas correctly. The SDL_BlitSurface() function now performs correctly even if the standard C library does not implement the memcpy() function safely for operations on overlapping memory areas.

BZ#640682

When running an SDL application in window mode and using certain window managers (like Fluxbox), the left-button event was not reported to the application. This update fixes dispatching of notifications when a window is left and the parent window is not interested in grabbing or ungrabbing events to handle window focus changes. With this update, the event of pressing the mouse button is reported to the application even if the window manager handles grab events.

BZ#640694

When using a hat-type analog joystick in an SDL application, the application terminated unexpectedly with a segmentation fault when the user moved the hat. The data structure which defines the analog part of a joystick has been updated to match the data structure passed by the kernel to the SDL library. The state of the analog hat is now properly passed by the kernel to the SDL library and it is interpreted accurately by the library.

BZ#640687

Prior to this update, the SDL spec file contained two invalid configuration options. This update removes these unrecognized options from the spec file. All options defined in the spec file are now recognized by the SDL configure script.

All users of SDL are advised to upgrade to this updated package, which fixes these bugs.

5.296. SEABIOS

5.296.1. RHBA-2012:0802 – seabios bug fix and enhancement update

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

The seabios package contains a legacy BIOS implementation, which can be used as a coreboot payload.

Bug Fixes

BZ#757999

Previously, SeaBIOS sometimes booted from an incorrect drive. This happened because the QEMU

hard-drive priority was lower than the virtio block-device priority. With this update, the QEMU hard-drive priority has been raised above the virtio block-device priority and SeaBIOS now boots from the correct drive.

BZ#771946

Previously, a guest could remain unresponsive during boot after the S3 (Suspend to RAM) state as SeaBIOS failed to advertise to the guest's operating system that the device was powered down. With this update, the underlying code handling the block device resume has been fixed and the problem no longer occurs.

BZ#786142

Previously, a Windows guest could detect an HPET (High Precision Event Timer) device although the guest had the HPET device disabled. This occurred because the HPET device was defined in the DSDT (Differentiated System Description Table). This update removes the definition from the table and the problem no longer occurs.

BZ#801293

Booting from some USB flash drives could fail because SeaBIOS did not support recovery from USB STALL conditions. This update adds support for recovery from STALLs.

BZ#804933

RTC (Real-Time Clock) wake-up for Windows guest did not work. With this update, the underlying code of FADT (Fixed ACPI Description Table) has been fixed to match QEMU behavior and the problem no longer occurs.

BZ#808033

Previously, if a device was hot plugged while the guest was still processing a previous hot-plug event, the new hot-plug event failed to be processed and the device was not detected. With this update, SeaBIOS uses a different event to handle hotplugging and the problem no longer occurs.

BZ#810471

Guest booting could fail if the guest had more than 62 sockets and multiple virtio disk devices. This happened because, BIOS ran out of memory and failed to initialize the boot disk. With this update, new memory is allocated under these circumstances and booting succeeds.

Enhancements

BZ#809797

The in-guest S4 (Suspend-to-Disk) and S3 (Suspend-to-RAM) power management features were added as a Technology Preview. The features provide the ability to perform suspend-to-disk and suspend-to-RAM functions on the guest. To enable the feature, users have to choose the /usr/share/seabios/bios-pm.bin file for VM BIOS instead of the default /usr/share/seabios/bios.bin file through libvirt.

BZ#782028

SeaBIOS now supports booting from virtio-scsi devices.

More information about Red Hat Technology Previews is available here:

https://access.redhat.com/support/offerings/techpreview/

All seabios users are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.297. SED

5.297.1. RHBA-2012:0955 – sed bug fix update

Updated sed packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

Sed is a stream or batch non-interactive editor. It takes text as input, performs an operation or set of operations on the text, and outputs the modified text.

Bug Fixes

BZ#709956

When the "--copy" option was used with the sed utility, sed created a back-up copy of the original file. However, the back-up file was not deleted after a successful operation. A patch has been provided to address this bug and back-up files are now properly deleted in the described scenario.

BZ#724962

Previously, the sed test suite was setting locale for UTF-8 variants incorrectly. Consequently, operation of the rpmbuild utility could fail, among other issues. This bug has been fixed and the test suite now sets the locale correctly.

BZ#812316

A static analysis tool discovered several minor file I/O resource leaks, which occurred when sed encountered an error. This bug has been fixed and the leaks no longer occur.

All users of sed are advised to upgrade to these updated packages, which fix these bugs.

5.298. SELINUX-POLICY

5.298.1. RHBA-2012:1581 – selinux-policy bug fix update

Updated selinux-policy packages that fix the bug 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#878360

Due to a bug in the SELinux policy, it was not possible to run a cron job with a valid MLS (Multi Level Security) context for the sysadm_u SELinux user. This update fixes relevant SELinux policy rules and cron now works as expected in the described scenario.

BZ#886210

Previously, SELinux prevented "rhevm-guest-agent-gdm-plugin" to connect to the SO_PASSCRED UNIX domain socket. Consequently, Single Sign-On (SSO) did not work because the access to the credential socket was blocked. This update fixes the relevant policy and SSO now works as expected in the described scenario.

All users of SELinux are advised to upgrade to these updated packages, which fix this bug.

5.298.2. RHBA-2012:1441 – selinux-policy bug fix update

Updated selinux-policy packages that fix two bugs 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#864366

Previously, SELinux was blocking the /usr/libexec/qemu-kvm utility during a migration of a virtual machine from Red Hat Enterprise Virtualization Manager. Consequently, such a migration attempt failed and AVC messages were returned. This update fixes the virt_use_fusefs boolean and adds the sanlock_use_fusefs boolean, thus allowing the migration to succeed in the described scenario.

BZ#867395

When trying to start a virtual machine on a POSIX-compliant file system, SELinux denied the operation and returned AVC messages. This update amends the SELinux policy to allow the described scenario to succeed.

Users of selinux-policy are advised to upgrade to these updated packages, which fix these bugs.

5.298.3. RHBA-2013:0002 – selinux-policy bug fix update

Updated selinux-policy packages that fix one bug 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 Fix

BZ#888381

Previously, the quota_db type was created as the openshift_var_lib_t type. Consequently, an attempt to create a quota system on openshift_var_lib_t failed with a permission error. The relevant part of the SELinux policy has been fixed and the quota system can now be created as expected.

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

5.298.4. RHBA-2012:1252 – selinux-policy bug fix update

Updated selinux-policy packages that fix two bugs 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#840674

Previously, with the MLS policy activated, a user created with a MLS level was not able to log into the system using the ssh utility because an appropriate MLS policy rule was missing. This update adds the MLS rule and users can now log into the system as expected in the described scenario.

BZ#852456

When OpenMPI (Open Message Passing Interface) was configured to use the parallel universe environment in the Condor server, a large number of AVC messages was returned when an OpenMPI job was submitted. Consequently, the job failed. This update fixes the appropriate SELinux policy and OpenMPI jobs now pass successfully and no longer cause AVC messages to be returned.

Users of selinux-policy are advised to upgrade to these updated packages, which fix these bugs.

5.298.5. RHEA-2012:1471 – selinux-policy enhancement update

Updated selinux-policy packages that add an enhancement 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.

Enhancement

BZ#876075

An SELinux policy for openshift packages has been added.

Users of selinux-policy are advised to upgrade to these updated packages, which add this enhancement.

5.298.6. RHBA-2012:1004 – selinux-policy bug fix update

Updated selinux-policy packages that fix one bug 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 Fix

BZ#833053

When the system produces a new SELinux denial, the setroubleshootd daemon executes the rpm tool to check information about the relevant packages. Previously, setroubleshootd was unable to execute the rpm tool, and AVC denials were logged in the /var/log/audit/audit.log file. With this update, the relevant policy has been corrected so that SELinux denials are no longer produced in the described scenario.

All users of selinux-policy are advised to upgrade to these updated packages, which fix this bug.

5.298.7. RHBA-2012:0780 – selinux-policy bug fix and enhancement update

Updated selinux-policy packages that fix a number of bugs 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#666332

Previously, the **sshd** init script tried to regenerate new keys during the **sshd** service startup and the **ssh-keygen** command failed to write public keys because of an incorrect SELinux security context

for the **ssh_host_rsa_key.pub** file. The security context has been updated and now the **sshd** service can start up correctly.

BZ#739886

Due to an error in an SELinux policy, SELinux incorrectly prevented the **rndc** service from reading the /**proc**/**loadavg** file. This update provides updated SELinux rules that allow **rndc** to read the /**proc**/**loadavg** file.

BZ#746961

When a non-root user (in the **unconfined_t** domain) ran the **ssh-keygen** utility, the SELinux policy did not allow **ssh-keygen** to create a key outside of the **~/.ssh** directory. This update adapts the relevant SELinux policy to make sure a key can be created by a non-root user in the described scenario.

BZ#748190

Previously, when a user tried to use the **selinux_avcstat** Munin plug-in, this caused Access Vector Cache (AVC) messages to be written to the audit log. With this update, a new SELinux policy has been provided for **selinux_avcstat** to fix this bug.

BZ#748971

Due to an incorrect SELinux policy, SELinux prevented the **openswan** utility to use the labeled IPsec protocol. This update provides updated SELinux rules and allows **openswan** to label IPsec as expected.

BZ#749311

Previously, the **nagios** event handlers were not supported by any SELinux policy, which broke their functionality. With this update, this support has been added to SELinux policy and **nagios** event handlers now work correctly with SELinux.

BZ#749501

Previously, when SELinux was running in Enforcing mode, the **google-chrome** program was unable to execute the **nacl_helper_bootstrap** command. This update provides an updated SELinux security context and rules that allow **google-chrome** to execute **nacl_helper_bootstrap**.

BZ#750869

Previously, the SELinux Multi-Level Security (MLS) policy did not allow users to use either the **newrole** or **sudo** command together with the **sssd** service configured, when the user was logged in the **wuth** custom MLS range. This update fixes the relevant SELinux policy to allow users to use this configuration.

BZ#751558

With SELinux in Enforcing mode, running the **mail** program as root with the **unconfined.pp** policy module disabled resulted in a permission to be denied and an AVC message to be generated. This update fixes relevant SELinux policy rules to allow the **mail** program to run properly in the described scenario.

BZ#751732

Due to an error in an SELinux policy, SELinux incorrectly prevented the **subscription-manager** service from reading the /**proc**/**2038**/**net**/**psched** file. This update provides updated SELinux rules that allow **subscription-manager** to read that file.

BZ#752418

Prior to this update, the **pyzor** application was denied the permission to write to the **ABRT** socket file. Consequently, an AVC message was reported. This update corrects the SELinux policy to grant **pyzor** the necessary permission in the described scenario.

BZ#752924

With SELinux running in Enforcing mode, the **smbcontrol** program was unable to send a signal to itself. Consequently, AVC messages were written to the audit log. This update fixes the relevant policy to support this operation.

BZ#718273

Previously, when SELinux was running in Enforcing mode, **gridengine mpi** jobs were not started correctly. A new policy for these jobs has been provided and **gridengine mpi** jobs now work as expected.

BZ#753184, BZ#756498

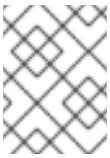
Previously, user **cron** jobs were set to run in the **cronjob_t** domain when the SELinux MLS policy was enabled. As a consequence, users could not run their **cron** jobs. The relevant policy rules have been modified and user **cron** jobs now run in the **user** domain, thus fixing this bug.

BZ#753396

When running the **libvirt** commands, such as **virsh iface-start** or **virsh iface-destroy**, with SELinux in Enforcing mode and NetworkManager enabled, the commands took an excessive amount of time to finish successfully. With this update, the relevant policy has been added and **libvirt** commands now work as expected.

BZ#754157

When the **auditd** daemon was listening on port 60, the SELinux Multi-Level Security (MLS) policy prevented **auditd** from sending audit events to itself from the same system if it was also running on port 61. This update fixes the relevant policy and this configuration now works as expected.



NOTE

Before the fix, the described scenario was possible to perform with the use of the **audisp-remote** plug-in.

BZ#754455

With SELinux enabled, the **rsyslogd** daemon was unable to start because it was not previously allowed to run the **setsched** operation using the Transport Layer Security (TLS) protocol. This update corrects the relevant SELinux policy and **rsyslogd** now starts as expected.

BZ#755877

With SELinux in Enforcing mode, the **ssh-keygen** utility could not access various applications and thus could not be used to generate SSH keys for such applications. With this update, the **ssh_keygen_t** SELinux domain type has been implemented as unconfined, which ensures the **ssh-keygen** utility works correctly.

BZ#759403

The **ssh-keygen** utility was not able to read from and write to the **/var/lib/condor/** directory. Consequently, with SELinux in Enforcing mode, an OpenMPI job submitted to the parallel universe

environment failed to generate SSH keys. With this update, a new SELinux policy has been provided for the **/var/lib/condor/** directory, which allows **ssh-keygen** to access this directory as expected.

BZ#759514

When running a KDE session on a virtual machine with SELinux in Enforcing mode, the session was not locked as expected when the SPICE console was closed. This update adds necessary SELinux rules, which ensure that the session is properly locked in the described scenario.

BZ#760537

Previously, the **/var/www/vweb1/logs/** directory was labeled as **httpd_log_t**, which blocked access to parts of additional web space. With this update, the **httpd_log_t** security context has been removed for this directory, thus fixing this bug.

BZ#767195

With SELinux in Enforcing mode, the **httpd** service could not read Git files with the **git_system_content_t** security label. This update corrects the relevant SELinux policy rules to allow **httpd** to read these Git files.

BZ#767579

Due to an error in an SELinux policy, SELinux incorrectly prevented to set up a quota on a file system, which was mounted as an user home directory, if the **quotacheck -c /user/home/directory** command was used. This update provides updated SELinux rules that allow to properly set up quotas in the described scenario.

BZ#754646

Previously, SELinux prevented the **sanlock** daemon from searching NFS directories. This update provides the **sanlock_use_nfs** boolean variable to fix this bug.

BZ#768065

When running the Postfix email server, the Amavis virus scanner, and the Spamassassin mail filter on Red Hat Enterprise Linux 6, the **spamc_exec_t** and **razor_exec_t** files were alias files, thus referencing the same context. Consequently, the **restorecon** utility reported these mislabeled files as related to the **razor** application. With this update, the **razor.pp** policy file has been removed and **restorecon** no longer reports these mislabeled files.

BZ#769301

Previously, if SSSD (System Security Services Daemon) used the **keyctl_join_session_keyring()** and **keyctl_setperm()** functions to connect to the kernel keyring and store passwords securely while the **sssd** daemon was running, it was permitted by SELinux. This update fixes the relevant SELinux policy rules to allow the SSSD **sys_admin** capability to process these operations properly.

BZ#769352

An incorrect SELinux policy prevented the **qpidd** service from starting. This update provides updated SELinux rules, which allow **qpidd** to be started correctly.

BZ#769819

Due to the labeling change for the **/var/spool/postfix/deferred** directory, the Postfix email server terminated. This update provides updated SELinux rules to allows Postfix to run as expected.

BZ#769859

Previously, when installing an updated selinux-policy-targeted package on a system with SELinux disabled, the following error messages were returned:

```
SELinux: Could not downgrade policy file /etc/selinux/targeted/policy/policy.24, searching for an older version.
SELinux: Could not open policy file -- /etc/selinux/targeted/policy/policy.24:  No such file or directory
load_policy: Can't load policy:  No such file or directory
```

This update provides the updated SELinux spec file that tests SELinux status correctly in the described scenario, thus preventing this bug.

BZ#773641

When SELinux was running in Enforcing mode, the **ssh-keygen** utility was unable to write to NFS home directories due to missing SELinux policy rules. This update provides updated SELinux rules that allow **ssh-keygen** to write to NFS home directories using the **use_nfs_home_dirs** boolean variable.

BZ#782325

When the user tried to execute the **check_disk** Munin plug-in on a remote system via NRPE (Nagios Remote Plugin Executor), the permission was denied and an AVC message was generated. This update fixes relevant SELinux policy rules to allow **check_disk** to read the **/sys/** directory, thus fixing this bug.

BZ#783592

Previously, SELinux policy for the **ipa_memcached** service was missing. Consequently, **ipa_memcached** did not work correctly with SELinux in Enforcing mode. This update adds support for **ipa_memcached**, thus fixing this bug.

BZ#784011

With the MLS SELinux policy enabled, an administrator running in the **sysadm_t** SELinux domain was not able to run the **rpm** command. This update provides updated SELinux rules to allow administrators to run **rpm** in the described scenario.

BZ#786597

Previously, when SELinux was running in Enforcing mode, the mail-related Munin plug-ins were not able to access the **/var/lib/** directory. Consequently, these plug-ins could not work correctly. This update provides updated SELinux rules, which allow these plug-ins to access **/var/lib/** and work as expected.

BZ#787271

If a custom cluster MIB (Management Information Base) implementation was run as a separate process, SELinux in Enforcing mode prevented the **snmpd** service to connect through the AgentX (Agent Extensibility) protocol. This bug has been fixed and the updated SELinux policy rules now allow to run custom cluster MIB implemantions.

BZ#788601

With SELinux in Enforcing mode, the **httpd** service was unable to access link files in the **/var/lib/zarafa/** directory, which caused various problems for the Zarafa groupware with DRBD (Distributed Replicated Block Device) support. This update provides updated SELinux rules and allows **httpd** to access the directory and Zarafa now works as expected.

BZ#788658

With SELinux in Enforcing mode, an OpenMPI job submitted to the parallel universe environment failed on SSH key generation. This happened because the **ssh-keygen** utility was unable to access the **/var/lib/condor/** directory. This update provides a new SELinux policy for **/var/lib/condor/**, which allows **ssh-keygen** to read from and write to this directory, thus fixing this bug.

BZ#789063

With SELinux in Enforcing mode, restarting the **tgtd** service resulted in SELinux AVC denial messages being returned when **tgtd** was not able to read the **abi_version** value. This update fixes the relevant SELinux policy rules to allow **tgtd** to read **abi_version**.

BZ#790980

If a custom home directory was set up as an NFS home directory, the **google-chrome** application was not able to write to this home directroy. With this update, the **use_nfs_home_dirs** variable has been fixed and **google-chrome** can now write to the NFS home directory in the described scenario.

BZ#791294

An incorrect SELinux policy prevented the **qpidd** service from connecting to the AMQP (Advanced Message Queuing Protocol) port when the **qpidd** daemon was configured with Corosync clustering. This update provides updated SELinux rules, which allow **qpidd** to be started correctly.

BZ#796351

Previously, SELinux received AVC denial messages if the **dirsrv** utility executed the **modutil -dbdir /etc/dirsrv/slapd-instname -fips** command to enable FIPS mode in an NSS (Network Security Service) key/certificate database. This happened because the **NSS_Initialize()** function attempted to use pre-link with the **dirsrv_t** context. With this update, the pre-link is allowed to re-label its own temporary files under these circumstances and the problem no longer occurs.

BZ#799102

With SELinux in Enforcing mode, Samba could not connect to dirsrv/slapd (389DS) via LDAPI, which caused AVC denial messages to be returned. Also, the **dirsrv** service failed to start properly due to this issue. This update provides an updated SELinux context for the **/var/run/slapd.*** socket and these services can be started as expected now.

BZ#799968

SSSD sometimes handles high load systems with more than 4,000 processes running simultaneously. Previously, SELinux in Enforcing mode produced an AVC message related to the **CAP_SYS_RESOURCE** privilege, which is needed to request a higher open file-descriptor limit. With this update, a new SELinux policy rule has been added to allow the **CAP_SYS_RESOURCE** capability for the SSSD service.

BZ#801163

With SELinux in Enforcing mode, the **chsh** utility did not work on servers that authenticated with Kerberos. SELinux prevented **chsh** from accessing certain files and directories. Now, updated SELinux rules have been provided to allow **chsh** to work properly in the described scenario.

BZ#802247

When a directory was mounted using NFS, restarting the **nfsclock** service produced an AVC denial message then reported to the **/var/log/audit/audit.log** log file. Updated SELinux policy rules have been provided, which allow the **rpc.statd** binary to execute the **sm-notify** binary, and restarting **nfsclock** now works properly.

BZ#802745

When files were created by the **/usr/bin/R** utility in user home directories, an incorrect SELinux context type of **user_home_dir_t** was returned, rather than the expected **user_home_t** context. This update fixes the relevant SELinux policy rules to allow **/usr/bin/R** to create directories in user home directories with correct labeling.

BZ#803422

When an ext4 partition was mounted using NFS, running the **xfstest** utility on this partition failed because write operations were denied on this partition. With this update, appropriate SELinux policy rules have been provided and write operations are now allowed to such partitions in the described scenario.

BZ#804024

Previously, installation of the selinux-policy-minimum package failed because a scriptlet of this policy attempted to access the **/etc/selinux/targeted/seusers** file. Now, the **selinux-policy.spec** file has been modified to store its users' information separately and selinux-policy-minimum can be installed properly.

BZ#804186

Previously, the Postfix email server was unable to work properly with the **~/Maildir/** set up. To fix this bug, a new SELinux context has been provided for the **/root/Maildir/** directory.

BZ#804922

With SELinux enabled, a Red Hat Enterprise Linux 6.2 client, which queried an NFS server also running on Red Hat Enterprise Linux 6.2, to get quota details, resulted in no output on the client and the following message to be reported to the server's logs:

> rpc.rquotad: Cannot open quotafile aquota.user and the associated AVC.

Updated SELinux policy rules, which allow this type of queries between NFS client and server, have been provided, thus fixing this bug.

BZ#805217

Previously, with SELinux in Enforcing mode and the **internal-sftp** subsystem configured, users with the **unconfined_t** SELinux type were unable to connect using the **sftp** utility. This update fixes the SELinux policy to allow users to utilize **sftp** successfully in the described scenario.

BZ#807173, BZ#820057

Due to the **nfs_export_*** booleans values being removed from Red Hat Enterprise Linux 6.3, users could not export subdirectories under the **/tmp/** directory and the mounting operations to such directories also failed. With this update, appropriate rules have been provided to allow users to perform these actions in the described scenario.

BZ#807456

With SELinux in Enforcing mode, the **cgconfig** service could not be started if an NIS (Network Information Service) user was specified in the **/etc/cgconfig** file. This update fixes the relevant SELinux policy rules and allow **cgconfig** to use NIS properly.

BZ#808624

When the Dovecot LMTP (Local Mail Transfer Protocol) server was configured as a virtual delivery

agent on a Postfix-based mail server, the **sieve** script was not working correctly with SELinux in Enforcing mode. This update provides appropriate SELinux policy rules to allow the **sieve** script to work correctly in the described scenario.

BZ#809746

Due to an incorrect SELinux policy, the **heartbeat** service could not be started correctly. New SELinux policy rules have been provided to allow **heartbeat** to execute the **/usr/lib64/heartbeat/plugins/InterfaceMgr/generic.so** binary, thus fixing this bug.

BZ#812850

With SELinux in Enforcing mode, the **service libvirt-qmf restart** command caused AVC denial messages to be logged to the **/var/log/audit/audit.log** file. This update fixes the relevant SELinux policy rules and the command no longer produces AVC messages.

BZ#812854

Previously, the **package-cleanup** utility did not work properly when called from a **cron** job. To fix this bug, the **/usr/bin/package-cleanup** binary has been labeled with the **rpm_exec_t** SELinux policy label and **package-cleanup** now works as expected in the described scenario.

BZ#813803

Previously, the **system-config-kdump** utlity did not work properly with SELinux enabled. To fix this bug, the **/etc/zipl.conf** file has been labeled with the **boot_t** SELinux security label.

BZ#814091

Fence agents (of the fence-agents package) in Red Hat Cluster Suite can use several different methods to connect to fencing devices. While using **telnet** or **ssh** works correctly under SELinux, some agents use SNMP. However, the **snmpwalk**, **snmpget**, and **snmpset** utilities did not work due to an incorrect SELinux policy. SELinux policy rules have been updated to allow SNMP utilities running with the **fenced_t** security type to be able to create files under the **/var/lib/net-snmp/** directory, thus fixing this bug.

BZ#821004

With the SELinux MLS policy enabled, the **sysadm_r** SELinux role could not create a cron job for another user. This bug has been fixed and the **sysadm_r** SELinux role now belongs among cron admin roles, thus fixing this bug.

Enhancements

BZ#727145

A new policy for the **cfengine** service has been added to make the system management work while using **cfengine**.

BZ#747239

This update provides a new SELinux policy for the **quota-nld** service.

BZ#747993

This update provides a new SELinux policy for the **flash** plug-in. Previously, the **plugin-container** processes of this plug-in were running as unconfined.

BZ#749200

This update provides new SELinux policies for the **matahari-qmf-sysconfigd** and **matahari-qmf-sysconfig-consoled** services.

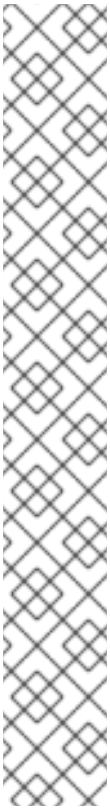
BZ#760405

The following boolean variables have been removed because they no longer had any effect:

```
allow_nfsd_anon_write
nfs_export_all_rw
nfs_export_all_ro
```

BZ#787413

Previously, there was no separation between the **secadm_r**, **sysadm_r** and **auditadm_r** SELinux roles related to certain operations with log files. This update introduces the new **sysadm_secadm.pp** SELinux module to provide the role separation.



NOTE

Note that if the **sysadm_secadm.pp** module is disabled, **sysadm_r** is unable to modify security files in the **/var/log/** directory, which only **secadm_r** can do. The basic separation of the roles is as follows:

- The **auditadm_r** role is able to modify the **/var/log/audit.log** log file.
- The **secadm_r** role is able to modify various SELinux properties as well as files in the **/var/log/** directory with necessary level. Users of this role can also change a level or a SELinux state, or can load a new module.
- The **sysadm_r** role (with **sysadm_secadm** disabled) is able to modify all non-security files because **sysadm_r** is based on the **userdom_admin_user_template()** function, which contains the following directives:

```
files_manage_non_security_dirs($1_t)
files_manage_non_security_files($1_t)
```

Users of this role are *not* able to modify **/var/log/audit/audit.log**, the **auditd** daemon configuration files, or change a level or a SELinux state.

BZ#795474

Previously, the **rsync** utility could not access files in either NFS or CIFS home directories. The new **rsync_use_nfs** boolean value has been provided to provide support for both file systems.

BZ#798534, BZ#812932, BZ#818082, BZ#818611

Previously, the **privsep** parent process always ran in the **sshd_t** domain. Consequently, the **sshd_t** domain had to be relaxed more than necessary for user SSH processes. This update introduces new SELinux policy rules to support permission separation for user SSH processes, each of which now runs in user context as expected.

BZ#801015

A new SELinux policy support has been added for the **matahari-qmf-rpcd** service.

BZ#801408

With this update, over 400 man pages documenting all confined domains and users on the system have been provided. You can acccess them using commands such as the following:

```
man httpd_selinux
man staff_selinux
```

BZ#807682

This update adds SELinux support for **ssh_to_job** for VM/Java/Sched/Local universe.

BZ#807824

This update adds SELinux support for the Cherokee web server.

BZ#809356

This update adds a new SELinux policy for the **libvirt-qmf** service.

BZ#810273

This update adds SELinux support for the **lvmetad** daemon.

BZ#811532

With this update, support for extended file attributes (xattr) has been added for the **ZFS** file system.

BZ#821038

This update adds a new SELinux policy for all OpenStack services.

Users of selinux-policy should upgrade to these updated packages, which fix these bugs and add these enhancements.

5.298.8. RHBA-2013:0904 – selinux-policy bug fix update

Updated selinux-policy packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The selinux-policy packages contain the rules that govern how confined processes run on the system.

Bug Fix

BZ#966996

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.

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

5.299. SERVICELOG

5.299.1. RHBA-2012:0989 – servicelog bug fix update

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

The servicelog packages provide a relational database to help manage errors on IBM eServer System p machines. Firmware and device driver errors are logged and managed using this database, which provides advanced error management capabilities.

Bug Fix

BZ#814160

Prior to this update, the "servicelog_manage -help" option returned a Null value instead of a help message. Also, repair actions were not properly deleted. This update modifies the underlying code so that a more informative help message is returned and repair actions are deleted along with regular events.

All users of servicelog are advised to upgrade to these updated packages, which fix this bug.

5.300. SETROUBLESHOOT

5.300.1. RHBA-2012:1005 – setroubleshoot bug fix update

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

The setroubleshoot packages provide tools to help diagnose SELinux problems. When AVC messages are returned, an alert can be generated that provides information about the problem and helps to track its resolution. Alerts can be customized to user preference. The same tools can used to process log files.

Bug Fix

BZ#832186

Previously, SELinux Alert Browser did not display alerts even if selinux denials were presented. This was caused by sedispatch, which did not handle audit messages correctly, and users were not able to fix their SELinux issues according to the SELinux alerts. With this update, this bug has been fixed so that users are now able to fix their SELinux issues using setroubleshoot, as expected.

All users of setroubleshoot are advised to upgrade to these updated packages, which fix this bug.

5.300.2. RHBA-2012:0781 – setroubleshoot bug fix update

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

The setroubleshoot packages provide tools to help diagnose SELinux problems. When AVC messages are returned, an alert can be generated that provides information about the problem and helps to track its resolution. Alerts can be customized to user preference. The same tools can be used to process log files.

Bug Fixes

BZ#757857

Previously, when the system produced a large amount of SELinux denials, the setroubleshootd daemon continued to process all of these denials and increased its memory consumption. Consequently, a memory leak could occur. This bug has been fixed and memory leaks no longer occur

in the described scenario.

BZ#633213

When a duplicated command-line option was passed to the sealert utility, sealert terminated with an error message. With this update, the duplicate options are ignored and sealert works as expected, thus fixing this bug.

BZ#575686

Previously, some translations for setroubleshoot were incomplete or missing. This update provides complete translations for 22 languages.

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

5.301. SETUP

5.301.1. RHBA-2012:1367 – setup bug fix update

Updated setup packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The setup packages provide a set of important system configuration and setup files, such as passwd, group, and profile.

Bug Fixes

BZ#791140

Prior to this update, the "/etc/profile" script used a non-portable method for undefining the pathmunge() function. As a consequence, the script could encounter problems when using the korn shell (ksh). This update modifies the undefining method of the function to work more efficiently with alternative shells.

BZ#839410, BZ#860221

Prior to this update, the accounts for the haproxy system user, the jbosson-agentsystem user, and the jbosson system group were created with dynamic uid/gid assignment, which is not recommended for network daemons and for sensitive data. With this update, the static uid/gid pair 188:188 can be used to create these users and groups.

All users of setup are advised to upgrade to these updated packages, which fix these bugs.

5.301.2. RHBA-2012:0778 – setup bug fix and enhancement update

An updated setup package that fixes three bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

The setup package contains a set of important system configuration and setup files, such as passwd, group, and profile.

Bug Fixes

BZ#771388

Prior to this update, the /etc/filesystems configuration file did not contain a line with the ext4 file

system. This could lead to various problems; for example, a process that used the file to determine supported file systems was not able to recognize ext4 as a valid file system. This update adds the missing line in the /etc/filesystems file.

BZ#710185

Prior to this update, the /etc/services configuration file contained an entry with the Internet Assigned Numbers Authority (IANA) reservation of port 0 for the spr-itunes service. However, the reservation of port 0 does not represent a real port reservation (it is only acknowledgment of IANA that the service exists). The spr-itunes entry has been commented out in the /etc/services file and an extended comment has been added to clarify the issue.

BZ#724007

Prior to this update, the /etc/group configuration file contained unnecessary supplementary groups - especially the root groups posed some potential security risk. These groups were legacy remnants and are no longer required. To mitigate the risk of making some future exploit more severe only because of the root's supplementary groups, the groups have been removed from the defaults.

Enhancements

BZ#772746

The wallaby package creates a user ID (UID) and a group ID (GID) pair, both with the name "wallaby" and number 181. Prior to this update, the UID and GID pairs were not reserved by the setup package. As a consequence, other packages or system administrators could accidentally assign the values to other users and groups. With this update, the setup package reserves these UID/GID names and numbers, so that accidental UID/GID usage risk is reduced.

BZ#760178

The tog-pegasus-libs package creates a user ID (UID) and a group ID (GID) pair, both with the name "cimsrvr" and number 134. Prior to this update, the UID and GID pairs were not reserved by the setup package. As a consequence, other packages or system administrators could accidentally assign the values to other users and groups. With this update, the setup package reserves these UID/GID names and numbers, so that accidental UID/GID usage risk is reduced.

BZ#738294

The sanlock package creates a user ID (UID) and a group ID (GID) pair, both with the name "sanlock" and number 179. Prior to this update, the UID and GID pairs were not reserved by the setup package. As a consequence, other packages or system administrators could accidentally assign the values to other users and groups. With this update, the setup package reserves these UID/GID names and numbers, so that accidental UID/GID usage risk is reduced.

BZ#738177

The dhcp package creates a user ID (UID) and a group ID (GID) pair, both with the name "dhcpd" and number 177. Prior to this update, the UID and GID pairs were not reserved by the setup packages. As a consequence, other packages or system administrators could accidentally assign the values to other users and groups. With this update, the setup package reserves these UID/GID names and numbers, so that accidental UID/GID usage risk is reduced.

BZ#804203, BZ#804204, BZ#804205, BZ#806052

A new cloud engine feature requires new users and groups - namely aeolus, katello, elasticsearch and mongodb with numbers 180, 182, 183 and 184. Prior to this update, the UID and GID pairs were not reserved by the setup packages. To prevent accidental UID/GID usage by other packages or system

administrators, the aforementioned UID/GID names and number are now reserved by the setup package.

All users of setup are advised to upgrade to this updated package, which fixes these bugs and add these enhancements.

5.302. SLAPI-NIS

5.302.1. RHBA-2012:0821 – slapi-nis bug fix and enhancement update

Updated slapi-nis packages that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The slapi-nis packages contain the NIS server plug-in and the Schema Compatibility plug-in for use with the 389 directory server.

The slapi-nis packages have been upgraded to upstream version 0.40, which provides a number of bug fixes and enhancements over the previous version. (BZ#789152)

Bug Fixes

BZ#784119

Prior to this update, the schema compatibility plug-in could, under certain circumstances, leak memory when computing values for inclusion in the constructed entries even if the relevant values were not changed. As a consequence, the performance could decrease rapidly and all available memory was consumed. This update modifies the underlying code so that the memory leaks no longer occur.

BZ#800625

Prior to this update, the directory server could terminate unexpectedly when processing a distinguished name if the relative distinguished name of a compatibility entry contained an escaped special character. This update modifies the plug-in so that special characters are now escaped when generating relative distinguished name values.

BZ#809559

Prior to this update, padding values passed to %link were read as literal values. As a consequence, the values could not use the "%ifeq" expression. This update modifies the underlying code to treat the padding values as expressions using the "%ifeq" expression.

Enhancement

BZ#730434

Prior to this update, the plug-ins used the platform-neutral Netscape Portable Runtime (NSPR) read-write locking APIs to manage some of their internal data. This update modifies slapi-nis to use the locking functionality provided by the directory server itself.

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

5.303. SLF4J

5.303.1. RHBA-2012:1239 – slf4j bug fix update

Updated slf4j packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The Simple Logging Facade (SLF4J) for Java serves as a simple facade for various logging APIs allowing the end-user to plug in the desired implementation at deployment time.

Bug Fix

BZ#831933, BZ#828644

The slf4j packages contained a non-functional dummy API implementation which was not supposed to be used. This dummy implementation was always selected instead of other implementations and UnsupportedOperationException was thrown. The dummy API implementation has been removed, so that user-supplied implementation is now always chosen, and slf4j works as expected.

All users of slf4j are advised to upgrade to these updated packages, which fix this bug.

5.304. SMARTMONTOOLS

5.304.1. RHBA-2012:0803 – smartmontools bug fix and enhancement update

Updated smartmontools packages that fix various bugs and adds multiple enhancements are now available for Red Hat Enterprise Linux 6.

The smartmontools package contains two utility programs (smartctl and smartd) to control and monitor storage systems using the Self-Monitoring, Analysis and Reporting Technology System (SMART) built into most modern ATA and SCSI hard disks. In many cases, these utilities will provide advanced warning of disk degradation and failure.

The smartmontools package has been upgraded to version 5.42, which provides a number of bug fixes and enhancements over the previous version. This update also features improved support of SATA disks on 3ware 9750 RAID controllers, improved support of SSD devices, and a much larger database of supported devices. (BZ#697235, BZ#698317)

Bug Fix

BZ#784925

Prior to this update, the format in which certain HP SAS (Serial Attached SCSI) drives returned SMART data was not recognized by smartmontools and the following error message "scsi response fails sanity test" was logged. With this update, smartmontools now correctly recognizes the data and this error no longer occurs.

All users of smartmontools are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.305. SOS

5.305.1. RHSA-2012:0958 – sos bug fix and enhancement update

An updated sos package that fixes one security issue, several bugs and adds various enhancements is now available for Red Hat Enterprise Linux 6.

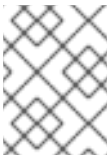
The Red Hat Security Response Team has rated this update as having low security impact. Common Vulnerability Scoring System (CVSS) base scores, which give detailed severity ratings, are available for each vulnerability from the CVE links associated with each description below.

The 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.

Security Fix

CVE-2012-2664

The sosreport utility collected the **Kickstart** configuration file (**/root/anaconda-ks.cfg**), but did not remove the root user's password from it before adding the file to the resulting archive of debugging information. An attacker able to access the archive could possibly use this flaw to obtain the root user's password. **/root/anaconda-ks.cfg** usually only contains a hash of the password, not the plain text password.



NOTE

This issue affected all installations, not only systems installed via **Kickstart**. A **/root/anaconda-ks.cfg** file is created by all installation types.

Bug Fixes

BZ#730641

Prior to this update, the path to the **/proc/net/** directory was specified incorrectly in SOS code. As a consequence, information necessary to debug certain bonding configurations from this directory was not available in the resulting archive. This update corrects the SOS networking module and ensures correct specification of the **/proc/net/** directory. As a result, generated sosreport tarballs contain the expected set of **/proc/net/** files.

BZ#749262

Previously, the **sosreport** utility failed to collect log files from Red Hat Network (RHN) Proxy Server installations. The problem was caused by an outdated package specification, which did not match the current package naming conventions. Consequently, logs that are sometimes required for RHN Proxy Server problem diagnostics were not collected automatically. This update corrects the package specification to match the current package naming conventions. As a result, RHN Proxy Server logs are collected correctly.

BZ#751273

Previously, output of the **brctl** command (used for Ethernet bridge configuration) was parsed incorrectly and caused **sosreport** to log errors. As a consequence, the **sosreport** command emitted a Python backtrace and certain bridge configuration information could not be collected. This update corrects the parsing of the **brctl** command output. As a result, no backtrace is emitted and all bridge configuration data is collected.

BZ#771393

Previously, SOS used inconsistent input sanitization rules. These rules varied depending on whether username and case information was supplied interactively, or was read from system configuration files. Consequently, SOS failed to properly sanitize certain invalid strings when read from configuration files, and applied different sanitization rules to the same strings when input

interactively. This update ensures that all name and number sanitization is carried out in a single location. As a result, name and number sanitization rules are now applied consistently, regardless of the source of the data.

BZ#782339

Previously, debug output produced by **sosreport** was limited due to changes to the logging subsystem introduced in SOS version 2.0. Consequently, very limited debug log information was collected as of that version of SOS. This update enhances the log subsystem and re-enables all previously disabled log messages. As a result, verbose log messages are now produced and recorded when requested via command-line options.

BZ#782589

Previously, **sosreport** did not correctly handle targets of symbolic links when copying files and directories into reports. Consequently, links in the report directory structure could have invalid targets. This update fixes the library routines dealing with file copying. The fix ensures that symbolic link targets are always copied when a requested path contains a symbolic link. As a result, **sosreport** handles symbolic link targets correctly and symbolic links in the report directory structure are always valid.

BZ#810702

Previously, SOS did not collect the machine check event (MCE) log from the **/var/log/mcelog** file. As a consequence, important information on the state of system hardware and previous hardware errors was sometimes missing in SOS reports. This update extends the SOS hardware module so that the MCE logs are collected when present in the **/var/log/mcelog** file. As a result, MCE log data is available in generated SOS reports.

BZ#812395

The IPA (Identity, Policy, Audit) identity and authentication components have been significantly updated in Red Hat Enterprise Linux 6.3. Consequently, the set of configuration and log data required to support these components has also changed. This update enhances the SOS IPA module and other related modules to collect information necessary for diagnosing problems in the new IPA versions. As a result, all information relevant for IPA diagnostics is collected from appropriately enabled systems running the updated IPA components.

BZ#814474

Previously, SOS used a single fixed path to collect all libvirt (virtualization API) logs from one directory. On some releases, the **libvirtd.log** file may be located in a different directory. Consequently, the **libvirtd.log** file was not collected on such systems. This update modifies **sosreport** so that it uses a wildcard matching both possible locations of the file. As a result, the **libvirtd.log** file is now collected on all supported releases.

Enhancements

BZ#739080

Previously, **sosreport** discarded program output from stderr (standard error stream). As a consequence, program warnings, diagnostics, and other messages were not included in reports generated by **sosreport**. This update modifies the way in which **sosreport** executes external programs. As a consequence, both stderr and stdout (standard output stream) messages returned by executed external programs are now included in reports generated by **sosreport**.

BZ#752549

Previously, SOS did not support the **GlusterFS** file system. As a consequence, running sosreport on a system where gluster packages were installed did not collect any Gluster-specific information from the system. This update adds a new plug-in that is necessary to collect the requisite logs for the Gluster product. As a result, information is collected from files located in the **/etc/glusterd/** and **/var/log/glusterfs/** directories. Several sets of command output are also collected to record the current state of the Gluster subsystem in the resulting report.

BZ#766583

Due to a previous update to the sos package, log files truncated for exceeding size limits are stored at a separate location in generated reports. This could be confusing for users unaware of this behavior. This update ensures that symbolic links to the truncated log files are added to the standard log file location. As a result, users and tools familiar with the standard location can now find truncated log files easily.

BZ#789096

Previously, the sos package contained a module for collection of general kernel information. The module did not collect additional information exposed by newer systems using the real-time kernel package (kernel-rt). This update adds a new kernel_real-time module and makes additions to the cgroups data collection. These changes result in collection of more complex diagnostic data on real-time kernel systems.

Users of sos should upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.306. SPICE-CLIENT

5.306.1. RHBA-2012:0888 – spice-client bug fix and enhancement update

Updated spice-client packages that fix several bugs and add multiple enhancements are now available for Red Hat Enterprise Linux 6.

The spice-client package provides the Simple Protocol for Independent Computing Environments (SPICE) client application. SPICE is a remote display protocol designed 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 KVM hypervisor or on Red Hat Enterprise Virtualization Hypervisors.

Bug Fix

BZ#552539

The SPICE client did not pass the volume multimedia keys events to the guest operating system. Therefore, it was not possible to change the guest's volume with these keys. Usage of the multimedia keys is now correctly caught by the client and passed to the guest system.

BZ#693431

In certain special multiscreen setups, after switching to full-screen mode and then quitting the SPICE client, the physical client screen turned off due to a bug in the code handling resolution switching. This updates fixes the code and the problem no longer occurs.

BZ#695323

The SPICE client did not properly support the Xinerama extension in full-screen mode on multi-

screen setups. Therefore, if the user switched to full-screen mode while using Xinerama, SPICE client windows failed to cover all physical screens used by the guest. This update improves Xinerama support and the SPICE client now behaves correctly in full-screen mode.

BZ#711810

Sound recording in the guest failed when another application was accessing the recording device on the SPICE client start-up. The client now uses the PulseAudio sound server, which allows multiple applications to access the recording device at the same time.

BZ#750030

The Red Hat Enterprise Virtualization console accessible from User Portal failed to open and returned error code 1032. This occurred on certain non-English locales as the value of the localized keyboard modifier was not considered a legal value of the hot-keys property. With this update, the hot-key value is parsed correctly even if unrecognized by the parser and falls back to its default value in such case.

BZ#791269

USB Auto-Share did not work on the initial full-screen of a SPICE session and a USB device could remain inaccessible unless the user switched focus to a different application and then back to the SPICE-client window. This occurred due to a race condition in the underlying code. The code has been modified and the problem no longer occurs.

BZ#791271

The SPICE context (right-click) menu was not always available in the client. The context menu was not displayed on clients where the USB Redirector Service was installed but not started and the user switched to window mode. The SPICE client has been updated and the context menu is now always available in window mode.

BZ#804561

Starting the SPICE client from the Red Hat Enterprise Virtualization portal failed when the local username was too long. With this update, the code has been modified and they are no longer constrains on the username length.

Enhancements

BZ#696075

Support for receiving a controller message telling the client to enable smart-card support has been added.

BZ#750856

The SPICE client is now able to handle requests from the Red Hat Enterprise Virtualization portal to enable or disable passing of the Ctrl+Alt+Delete key combination to the guest operating system.

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

5.307. SPICE-GTK

5.307.1. RHSA-2012:1284 – Moderate: spice-gtk security update

Updated spice-gtk packages that fix one security issue are now available for Red Hat Enterprise Linux 6.

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

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.

Security Fix

CVE-2012-4425

It was discovered that the spice-gtk setuid helper application, spice-client-glib-usb-acl-helper, did not clear the environment variables read by the libraries it uses. A local attacker could possibly use this flaw to escalate their privileges by setting specific environment variables before running the helper application.

Red Hat would like to thank Sebastian Krahmer of the SUSE Security Team for reporting this issue.

All users of spice-gtk are advised to upgrade to these updated packages, which contain a backported patch to correct this issue.

5.307.2. RHBA-2012:0767 – spice-gtk bug fix and enhancement update

Updated spice-gtk packages that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The spice-gtk packages provide a GTK2 widget for SPICE clients. Both virt-manager and virt-viewer can make use of this widget to access virtual machines using the SPICE protocol.

The spice-gtk packages have been updated to upstream version 0.11, which fixes multiple bugs and adds multiple enhancements. These packages also add support for native USB redirection. (BZ#773642)

Bug Fixes

BZ#772118

Prior to this update, the spice-gtk client could abort unexpectedly when the audio system failed to initialize. This update ignores audio initialization failures so that spice-gtk also works with failed audio system initialization.

BZ#805641

Prior to this update, memory leaks could, under certain circumstances, occur when the guest resolution was changed. This update modifies the underlying code so that changing the resolution no longer causes memory leaks.

BZ#807389

Prior to this update, videos showed a blue tint when using guests hosted on a Red Hat Enterprise Linux 5.8 host. This update uses the correct color conversion also on older SPICE servers so that videos are now correctly tinted.

BZ#809145

Prior to this update, copying large amounts of text could cause a stack overflow. As a consequence

SPICE clients that use the spice-gtk widget could abort with a segmentation fault. This update no longer allocates large amounts of data on the stack and large clipboard data can now be copied to the server.

Enhancement

BZ#758100

Prior to this update, the spice-gtk widget did not provide a native USB redirection. This update adds USB redirection support to spice-gtk.

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

5.308. SPICE-PROTOCOL

5.308.1. RHEA-2012:0760 – spice-protocol enhancement update

An updated spice-protocol package that adds several enhancements is now available for Red Hat Enterprise Linux 6.

The spice-protocol package contains header files that describe the SPICE protocol and the QXL para-virtualized graphics card. The SPICE protocol is needed to build newer versions of the spice-client and the spice-server packages.

BZ#758088

The spice-protocol package has been upgraded to upstream version 0.10.1, which provides a number of enhancements over the previous version, including support for USB redirection.

All users who build spice packages are advised to upgrade to this updated package, which adds these enhancements.

5.309. SPICE-SERVER

5.309.1. RHBA-2012:0765 – spice-server bug fix and enhancement update

Updated spice-server packages that fix multiple 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.

The spice-server packages have been upgraded to upstream version 0.10.1, which fixes multiple bugs and adds multiple enhancements. (BZ#758089)

Bug Fixes

BZ#741259

Prior to this update, the smart card channel looked for the error code at the wrong location. As a consequence, the error messages contained random code instead of the actual error code. This update modifies the smart card channel code so that correct error messages are now sent.

BZ#787669

Prior to this update, the server rejected connections without logging any information to the qemu log if the client provided a wrong password. This update modifies qemu-kvm so that the messages "Invalid password" or "Ticket has expired" are sent when the client provides a wrong password.

BZ#787678

Prior to this update, qemu did not log X.509 files. As a consequence, no output regarding certificates or keys was available. This update modifies the underlying code so that information on X.509 files is now available.

BZ#788444

Prior to this update, the "struct sockaddr" code in the spice server library API was too short to hold longer IPv6 addresses. As a consequence, the reported IPv6 address appeared to be broken or incomplete. This update modifies the underlying code to use "struct sockaddr_storate" that can now hold complete IPV6 addresses.

BZ#790749

Prior to this update, the default lifetime of the "SpiceChannelEventInfo" event was too short for the "main_dispatcher_handle_channel" event. As a consequence, freed memory could be accessed after the RedsStream was freed for the cursor and display channels. This update allocates the "SpiceChannelEventInfo" event together with allocating the "RedsStream" event, and deallocates it only after the "DESTROY" event.

BZ#813826

Prior to this update, the display driver could send bitmaps to the spice server that contained video frames, but were larger than the frames sent before. As a consequence, the larger frames were not synchronized with the video stream, and their display time could differ from the display time of other frames and the playback seemed to skip and interrupt. With this update, large bitmaps are directly attached to the video stream they contain. Now, the playback is smooth and no longer interrupts.

Enhancement

BZ#758091

Prior to this update, USB devices could not be redirected over the network. This update adds USB redirection support to the to spice-server.

All users requiring spice-server are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.310. SPICE-XPI

5.310.1. RHEA-2012:0956 – spice-xpi enhancement update

Enhanced spice-xpi packages are now available for Red Hat Enterprise Linux 6.

The spice-xpi package provides the Simple Protocol for Independent Computing Environments (SPICE) extension for Mozilla that allows the SPICE client to be used from a web browser.

The Simple Protocol for Independent Computing Environments (SPICE) is a remote display protocol designed for virtual environments. SPICE users can view a virtualized desktop or server from the local system or any system with network access to the server. SPICE is available for a variety of machine architectures and operating systems.

SPICE is used in Red Hat Enterprise Linux for viewing virtualized guests running on the KVM hypervisor or on Red Hat Enterprise Virtualization Hypervisors.

The spice-xpi packages have been upgraded to upstream version 2.7, which provides a number of enhancements over the previous version. In particular, the SPICE client started by the SPICE Firefox extension can now be configured through the update-alternatives mechanism. The currently available alternatives are spice-client and remote-viewer. (BZ#784846)

Enhancements

BZ#752090

The SPICE Firefox extension is now able to handle requests from the Red Hat Enterprise Virtualization portal to enable or disable passing the Ctrl+Alt+Delete key combination to the guest operating system and to pass those requests to the SPICE client.

BZ#641828

The SPICE Firefox extension is now able to handle requests from the Red Hat Enterprise Virtualization portal to enable or disable smartcard passthrough to the guest operating system and to pass those requests to the SPICE client.

BZ#807303

The SPICE Firefox extension is now able to handle requests from the Red Hat Enterprise Virtualization portal to enable or disable native USB redirection to the guest operating system and to pass those requests to the SPICE client.

BZ#747313

The SPICE Firefox extension is now able to handle requests from the Red Hat Enterprise Virtualization portal to enable or disable some graphical effects to the guest operating system and to pass those requests to the SPICE client. Disabling these effects can improve performance on WANs.

BZ#790416, BZ#823578

The SPICE Firefox extension is now able to correctly compose SSL channel names used for a SPICE session. Prior to this update, removing the prefix "s" was done only for channel names "main" and "inputs". The improvement consists of removing the prefix "s" from all the SSL channel names.

BZ#813231

Red Hat Enterprise Linux 6.3 includes an improved SPICE client called remote-viewer. This update changes the spice-xpi package to now require this improved client instead of requiring the spice-client package. As a result, when installing or upgrading spice-xpi, the virt-viewer package (containing the remote-viewer client) is installed instead of the spice-client package. The client that spice-xpi uses can be selected through the update-alternatives mechanism.

BZ#753155

The SPICE Firefox extension now provides log messages when spice-xpi functions are called from a

web page using JavaScript and for command line processes that spice-xpi runs (for example, spice-xpi-client). Whenever a web page calls the spice-xpi function, a DEBUG level message is written to the logs. It includes all variables that are passed to spice-xpi. By default, logs are written to the file "spice-xpi.log" in the "~/.spicec" directory; however, the log location can be changed by settings in the "logger.ini" file in the "/etc/spice/" directory.

Users of spice-xpi are advised to upgrade to these updated packages, which add these enhancements.

5.311. SQUID

5.311.1. RHBA-2012:1290 – squid bug fix update

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

[Updated 20th September 2012] This advisory has been updated with an accurate description of the "http10" option for BZ#852863. This update does not change the packages in any way.

Squid is a high-performance proxy caching server for web clients that supports FTP, Gopher, and HTTP data objects.

Bug Fixes

BZ#853053

Due to a bug in the ConnStateData::noteMoreBodySpaceAvailable() function, child processes of squid aborted upon encountering a failed assertion. An upstream patch has been provided to address this issue and squid child processes no longer abort in the described scenario.

BZ#852863

Due to an upstream patch, which renamed the HTTP header controlling persistent connections from "Proxy-Connection" to "Connection", the NTLM pass-through authentication does not work, thus preventing login. This update introduces the new "http10" option to the squid.conf file, which can be used to enable the change in the patch. This option is set to "off" by default. When set to "on", the NTLM pass-through authentication works properly, thus allowing login attempts to succeed.

BZ#852861

When the IPv6 protocol was disabled and squid tried to handle an HTTP GET request containing an IPv6 address, the squid child process terminated due to signal 6. This bug has been fixed and such requests are now handled as expected.

BZ#855330

The old "stale if hit" logic did not account for cases where the stored stale response became fresh due to a successful re-validation with the origin server. Consequently, incorrect warning messages were returned. With this update, squid no longer marks elements as stale in the described scenario, thus fixing this bug.

All users of squid are advised to upgrade to these updated packages, which fix these bugs.

5.312. SSSD

5.312.1. RHBA-2012:0747 – sssd bug fix and enhancement update

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

SSSD (System Security Services Daemon) provides a set of daemons to manage access to remote directories and authentication mechanisms. It provides NSS (Name Service Switch) and PAM (Pluggable Authentication Modules) interfaces toward the system and a pluggable back end system to connect to multiple different account sources.



NOTE

The sssd package has been upgraded to upstream version 1.8.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#735422)

Bug Fixes

BZ#818642

User authentication could fail if the user or its group data specified non-standard LDAP attributes due to incorrect handling of such attributes. With this update, such attributes are handled properly and user authentication now works under these circumstances as expected.

BZ#773655

Previously, SSSD did not correctly handle LDAP authentication requests failover under a heavy load and the request could fail with a system error. This occurred due to an invalid LDAP URI value if the second authentication request was sent before the first request could be processed by the failover service. With this update, the underlying code has been modified to ensure that the LDAP URI string remains valid until the LDAP authentication request is processed.

BZ#801407

The function handling pending requests on reconnect was checking an orphaned global variable that was never used. Consequently, if SSSD never received a response to a request, the request and any subsequent requests for the same information remained unhandled. With this update, the function refers correctly to the respective hash table and identical requests are now processed as expected even if the original request fails.

BZ#753842

SSSD uses libdbus for interprocess messaging. Previously, libdbus caused SSSD to terminate unexpectedly when SSSD passed it a username with a non-UTF-8 character. With this udpate, SSSD checks if the input contains non-UTF-8 characters and rejects requests with such characters gracefully.

BZ#822236

In the course of speeding up cached lookups for netgroups, SSSD inadvertently disabled the use of the nowait cache lookups. This functionality has now been restored and cache misses are reduced for oft-requested netgroups.

BZ#801451

When using IPA as access_provider, SSSD evaluated only HBAC (Host-Based Access Control Rules) rules and failed to evaluate password expiration policies in the PAM_ACCT_MGMT phase. Consequently, users who logged in to a FreeIPA-managed system with an alternative mechanism, such as SSH public-key or GSSAPI, did not have their password-expiration status evaluated and

managed to log in even if their accounts were expired or disabled. With this update, SSSD now checks password-expiration policies in the IPA access_provider and users with such accounts can no longer log in to the system in the scenario described.

BZ#787035

When looking up cached group entries, the glibc queries SSSD with a fixed buffer. If the group did not fit into the buffer, SSSD returned an error and glibc retried with an enlarged buffer. This caused performance issues when querying large groups as multiple retries involved repeated contacting of SSSD and reading the entries from the cache. The SSSD NSS client now keeps the group entry in memory until a sufficiently large buffer is provided and lookups for cached group entries are now faster.

BZ#803937

After an LDAP client-side migration, SSSD used the start TLS operation on a connection that was already encrypted by GSSAPI. Consequently, under certain circumstances, the sssd_be process, which communicated directly with the server, terminated unexpectedly and dumped core. With this update, the migration procedure has been fixed so that it now establishes a new TLS-only connection for the migration and the client-side password migration is more robust.

BZ#771706

Prior to this update, the SSSD daemon saved a NULL pointer instead of an empty service or a host group and later dereferenced the pointer if an IPA server contained an HBAC rule with these empty service groups or host groups. As a consequence, the NULL pointer dereference could abort SSSD. This update creates an empty array rather than using a NULL pointer. Now, SSSD handles empty service groups or host groups as expected.

BZ#742052

Prior to this update, SSSD performed a single LDAP search operation per every LDAP group member if the RFC2307bis schema was used. As a consequence, group lookups could take a long time especially for environments with large groups. This update leverages a "dereference" feature to allow downloading all the members in a single large search operation. Now, group lookups take significantly less time.

BZ#735827

The POSIX standard mandates that user and group names are case sensitive but the user and group names are case insensitive on Windows and on most LDAP servers. Name comparisons that matched in Windows did not match in Red Hat Enterprise Linux. This update introduces a new option, "case_sensitive", that allows to treat names in a case insensitive manner. This option is set to "true" by default, maintaining the POSIX standard setting.

BZ#735405

Prior to this update, the SSSD daemon printed a warning message to the /var/log/secure log if a user was passed to SSSD that SSSD could not handle, such as local users while processing logins for SSSD. As a consequence, the /var/log/secure log was filled with redundant error messages. With this update, the pam_sss.so module accepts the option "quiet" that suppresses the unknown user messages. Error messages about unknown users no longer appear in /var/log/secure.

BZ#766904

Prior to this update, SSSD only read its configuration at startup time and the verbosity of the debug logs could only be set at startup time. As a consequence, users had to leave noisy debug logs enabled for extended periods when trying to track down an intermittent error. A reboot to change

the debug level could cover the problem for some time until it reoccurred. With this update, a new command line tool is added to SSSD to change the debug level of live SSSD processes. Users can now change the debug verbosity of the SSSD processes without restarting SSSD.

BZ#785879

Prior to this update, configuration options were defined as required even when they were not required. Also the script for configuration parsing did not merge the old tree with the new one when changing certain options but created a new one and deleted the old one instead. As a consequence, the configuration file could change significantly, comments and blank lines disappeared, and also new options were added when updating a configuration file with scripts to parse the configuration. This update reduces the list of required options and modifies the configuration parsing script so it merges the old and the new tree. After being processed with the python scripts, the configuration file is now corresponding to the original.

BZ#785881

Prior to this update, the Identity Management provider used keytabs to authenticate against an Identity Management server by constructing the expected principal and then attempting to use this principal. If the constructed principal was not in the keytab, the entire operation failed and the backend was not able to connect to the Identity Management server. This update changes the approach to list all principals in a used keytab and to select the most convenient one. The current implementation uses a more flexible algorithm to find a suitable principal in a keytab.

BZ#785888

Prior to this update, the NSS responder used a negative cache to avoid asking repeatedly the provider for non-existent entities. The querying process for netgroups did not work efficiently with the negative cache. An empty netgroup could, under certain circumstances, be returned to the client even for a non-existent group. This update modifies the NSS responder to use a special flag indicating that the group was found in the cache when using a negative cache for netgroup lookups. Netgroup queries no longer return empty netgroups if they do not exist in the cache.

BZ#785902

Prior to this update, the SSSD cache storage function for user entities did not check empty strings in loginShell attributes. If the check encountered such an attribute, the storing procedure failed completely. When using a proxy provider and the utilized NSS module returned an empty loginShell, updating user records in the cache failed. This update ensures that the proxy provider does not pass empty strings to the function.

BZ#791208

Prior to this update, SSSD expected all users in a POSIX-enabled Active Directory group to be POSIX-enabled users. If some members of a POSIX-enabled group were lacking the POSIX username attribute, SSSD returned an error when looking up that group. This update ignores non-POSIX group members. SSSD now returns all POSIX-enabled group members and silently ignores non-POSIX members.

BZ#795562

Prior to this update, a server status in the SSSD server list was reset after 30 seconds to allow retries. If a full cycle over the server list took more than 30 seconds, the cycle started again. SSSD deployments using large server failover lists could loop indefinitely. This update modifies SSSD was to only loop over the fail over list once. If the SSSD tries all the servers in the fail over list without succeeding, the operation always fails.

BZ#798774

Prior to this update, SSSD used expand FQDN and DNS SRV to look up DNS SRV records for failover servers. On FreeIPA-enrolled machines, the client hostname could, under certain circumstances, not match the IPA domain name. These clients were unable to discover failover servers. When the id_provider is set to IPA, then the dns_discovery_domain is automatically set to the value of ipa_domain. FreeIPA clients are able to autodetect failover servers even if their hostname is not part of the FreeIPA domain.

BZ#799929

Prior to this update, SSSD was limited to using 1024 file descriptors for its sssd_nss and sssd_pam responder processes. On very busy systems with many user lookups and/or authentications, SSSD could run out of descriptors and stop responding to requests until it was restarted. This update increases the SSSD limit to 4096 descriptors. Users should not experience the resource exhaustion described above.

BZ#773660

Prior to this update, SSSD logged errors in the Kerberos authentication only into its own debug logs. Errors that occurred during Kerberos authentication are now sent to the syslog in addition to the debug logs.

BZ#772297

Prior to this update, the function for storing netgroups in SSSD cache did not check attributes that are contained in sysdb but not in the LDAP response from the server. If a netgroup has been cached by SSSD and it changed on the server in a way that it missed all the triples, this change was not projected in the cache. To avoid this problem, a check for attributes that are missing from the LDAP response when saving a netgroup has been added.

BZ#801533

Prior to this update, SSSD used a wrong counter and could access random memory when resolving a complex group structure during an initgroups operation. The random memory access terminated the sssd_be process. SSSD now uses the correct group counter and processes nested group structures correctly.

BZ#771702

In case SSSD was operating in the offline mode and a Kerberos password was requested with a configuration that also used the KDC server for changing passwords, SSSD was issuing the password change requests in an infinite loop. Specifically, the "sssd_be" process was looping infinitely and occasionally even terminating unexpectedly. The "sssd_be" process was fixed to not call the password changing request while operating in the offline mode. When a password change operation is requested while SSSD is offline, the operation exits gracefully.

BZ#805034

When an LDAP entry changed its attributes and was saved again into the SSSD cache, SSSD might have accessed an undefined variable value. This caused SSSD to crash. With this update, the variable is now initialized to a known default value, and SSSD no longer crashes when updating cached entries.

BZ#805108

Due to a programming error, a loop could only be exited when an error occurred. When a connection to a system with "knownhostproxy" enabled was closed, the loop was not exited and caused "sss_ssh_knownhostsproxy" to become unresponsive. This update fixes this bug so that the loop is exited when the connection is closed, and "sss_ssh_knownhostsproxy" no longer hangs.

BZ#768935

A bug in the SSSD configuration parser caused the parser library to terminate unexpectedly when an old SSSD configuration domain was removed and a new one was saved. Consequently, applications which used the configuration parser, such as "authconfig", would crash. This update fixes the SSSD configuration parser so that it no longer crashes.

BZ#814269

The OpenLDAP client libraries (used by SSSD) did not time out properly if communication with an LDAP server would drop packets instead of rejecting them. As a consequence, SSSD became unresponsive and never responded to requests. This update adds a timer to SSSD to ensure that connections are timed out after a reasonable amount of time; SSSD no longer hangs.

BZ#759186

When an SSSD service exited while a check for its presence was still in progress, SSSD might have accessed invalid memory, which resulted in a crash. With this update, any pending checks are canceled when an SSSD service exits, and SSSD no longer crashes.

BZ#746181

When a new group was added to the SSSD cache, it was not checked whether there was another group with the same GID already present in the database. With this update, when adding a new group to the cache, any group with the same GID that is already present in the cache is deleted.

Enhancements

BZ#761582

The **ldap_sasl_minssf** option has been added to the configuration of SSSD. This option can be used to specify the minimal level of encryption SSSD (or rather, the LDAP library used by SSSD) should use when communicating with a server.

BZ#739312

A new option, **ldap_chpass_update_last_change**, has been added to SSSD configuration. If this option is enabled, SSSD attempts to change the shadowLastChange LDAP attribute to the current time. Note that this is only related to a case when the LDAP password policy is used (usually taken care of by LDAP server), that is, the LDAP extended operation is used to change the password. Also note that the attribute has to be writable by the user who is changing the password.

BZ#742509

The **sss_cache** tool has been added to the SSSD package. This tool allows you to expire cached objects, which triggers their online renewal as soon as they are requested and it is possible to retrieve them from a server.

BZ#742510

SSSD had a single configurable option for setting the cache timeout for users, groups, netgroups and services. However, some deployments have different caching needs for different nsswitch maps. With this update, SSSD provides new options to configure each cache entry type's timeout individually:

```
entry_cache_user_timeout
entry_cache_group_timeout
entry_cache_netgroup_timeout
```

> entry_cache_service_timeout

Users can now define their cache timeouts on a per-entry basis. For more information about these options, refer to the sssd.conf(5) man page.

BZ#753763

SSSD has changed the behavior of the **debug_level** option in the "/etc/sssd/sssd.conf" file. For more information, refer to the Red Hat Enterprise Linux 6.3 Release Notes.

BZ#744197

SSSD now contains a configurable idle timeout, after which it disconnects from the LDAP server until the next request is received. As a result, SSSD is now a less resource-intensive client for LDAP servers.

BZ#805924

SSSD relies on some information it can retrieve from the RootDSE in order to determine the capabilities of the server. Some servers do not make the RootDSE available via unencrypted, non-authenticated LDAP bind (in violation of the LDAP standard). On such servers, SSSD operates in a slightly degraded mode, being unable to take advantage of any enhanced features of the LDAP server. With this update, SSSD now makes a second attempt to retrieve the RootDSE after it completes a successful bind attempt. SSSD is now able to take advantage of enhanced features on servers that do not expose the RootDSE to non-authenticated users.

BZ#728212

OpenLDAP servers sometimes report that paging control is available even if it is disabled. Consequently, SSSD previously attempted to use the paging control feature and failed to perform lookups relying on this feature, such as lookups of group members. With this update, a new option **ldap_disable_paging** has been added to SSSD, which allows the user to disable paging control on such servers manually.

BZ#736150

The ability to search multiple bases for each entry type has been added to SSSD.

BZ#761570

This update adds support for automount map caching as a Technology Preview. Cached automount maps allow a client machine to perform mount operations even though the LDAP server is unreachable. Also, the feature results in faster performance on the client and lower traffic on the LDAP server.

BZ#755506

To enable the behavior of **pam_check_host_attr**, users can now set the **ldap_access_order = host** and **ldap_user_authorized_host** options to enable access-control based on the presence of this attribute in LDAP.

BZ#766876

Evaluation of srchost HBAC rules can be unreliable and cause significant performance issues on login. With this update, SSSD now ignores srchost rules in HBAC processing by default. To enable the evaluation, set the newly-added **ipa_hbac_support_srchost** option to **true**.

BZ#753876

SSSD now supports querying the services map of LDAP and the proxy provider and users can have their services map served and cached.

BZ#766930

This update adds to SSSD the **override_homedir** option that allows the user to define a per-client override values for the home directory attribute.

BZ#785905

Prior to this update, SSSD debug messages provided precision down to the wallclock second. When debugging performance issues, users required higher precision in the timestamps. With this update, SSSD adds the "debug_microseconds" option to enable microsecond-level precision in debug messages. Users of SSSD now have the option to enable microsecond precision in debug log messages.

BZ#785907

A new option krb5_canonicalize has been added to SSSD configuration. When set to true, it sets a flag in krb5 request and the host and user principals are canonicalized and returned to SSSD by server. Note, that this feature requires Kerberos version 1.7 or later.

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

5.312.2. RHBA-2013:0677 – sssd bug fix update

Updated sssd packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

SSSD (System Security Services Daemon) provides a set of daemons to manage access to remote directories and authentication mechanisms. It provides NSS (Name Service Switch) and PAM (Pluggable Authentication Modules) interfaces toward the system and a pluggable back end system to connect to multiple different account sources.

Bug Fixes

BZ#847969

When the ldap_chpass_update_last_change option was enabled, the shadowLastChange attribute contained number of seconds instead of days. Consequently, when shadowLastChange was in use and the user was prompted to update their expiring password, shadowLastChange was not updated. The user then continued to get the error until they were locked out of the system. With this update, number of days is stored in shadowLastChange attribute and users are able to change their expiring passwords as expected.

BZ#867012

Kerberos options were loaded separately in the krb5 utility and the IPA provider with different codepaths. The code was fixed in krb5 but not in the IPA provider. Consequently, a Kerberos ticket was not renewed in time when IPA was used as an authentication provider. With this update, Kerberos options are loaded using a common API and Kerberos tickets are renewed as expected in the described scenario.

BZ#881460

When SSSD was built without sudo support, the ldap_sudo_search_base value was not set and the namingContexts LDAP attribute contained a zero-length string. Consequently, SSSD tried to set

ldap_sudo_search_base with this string and failed. Therefore, SSSD was unable to establish connection with LDAP server and switched to offline mode. With this update, SSSD considers the zero-length namingContexts value the same way as if no value was available, thus preventing this bug.

All users of sssd are advised to upgrade to these updated packages, which fix these bugs.

5.313. STRACE

5.313.1. RHBA-2012:1317 – strace bug fix and enhancement update

Updated strace packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The strace packages provide an utility to intercept and record the system calls called and received by a running process. The strace utility can print a record of each system call, its arguments, and its return value. The strace utility is useful for diagnosing, debugging, and instructional purposes.

Bug Fix

BZ#849052

Previously, the strace utility used magic breakpoints in the process startup code to detect and control process startup. Consequently, under certain circumstances, the %ebx register could be corrupted around the clone syscall within the libc_fork() function, which could cause an application to terminate unexpectedly with a segmentation fault while under strace control. This update changes strace to use the TRACE{FORK,VFORK,CLONE} ptrace capabilities which provide a cleaner, less error-prone interface to monitor and control process startup when tracing, thus preventing this bug.

All users of strace are advised to upgrade to these updated packages, which fix this bug.

5.314. SUBSCRIPTION-MANAGER

5.314.1. RHBA-2012:1073 – subscription-manager bug fix update

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

The Subscription Manager tool allows users to understand the specific products which have been installed on their machines, and the specific subscriptions which their machines are consuming.

Bug Fixes

BZ#833390

Prior to this update, CLI (command-line interface) calls to "subscription-manager release --list" were ignoring the command-line proxy options. As a consequence, running the command with proxy configuration values specified as arguments resulted in the proxy settings from the configuration file being used instead of those specified on the CLI. With this update, if the user specifies a proxy using command-line options, these values correctly override corresponding proxy settings in the rhsm.conf file when the command is run.

BZ#834558

Client ID certificates expire after one year, and previously could be regenerated only manually by the user. With this update, the client can automatically retrieve an updated client ID certificate from the entitlement server if this is supported by the target instance.

All users of subscription-manager are advised to upgrade to these updated packages, which fix these bugs.

5.314.2. RHBA-2012:1050 – subscription-manager bug fix update

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

The Subscription Manager tool allows users to understand the specific products which have been installed on their machines, and the specific subscriptions which their machines are consuming.

Bug Fixes

BZ#830267

Prior to this update, the Subscription list was blank in the Confirm Subscriptions page of Firstboot when LANG was set to ja_JP. This was because the rhn_review_gui screen and the rshm_confirm_subs screens had the same translated name in Japanese. This bug has been fixed and the Confirm Subscriptions page now correctly displays the Subscription list in the aforementioned case.

BZ#830269

When going back from the Subscription list in Firstboot, instead of being directed to the Entitlement Platform Registration page, the Set Up Software Updates page was displayed when LANG was set to ja_JP. This was because the rhn_review_gui screen and the rshm_confirm_subs screens had the same translated name in Japanese. This bug has been fixed and going back from the Subscription list page now works as expected in the scenario described.

All users of subscription-manager are advised to upgrade to these updated packages, which fix these bugs.

5.314.3. RHBA-2012:13347 – subscripton-manager bug fix update

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

The Subscription Manager tool allows users to understand the specific products which have been installed on their machines, and the specific subscriptions which their machines are consuming.

Bug Fixes

BZ#830267

Prior to this update, the Subscription list was blank in the Confirm Subscriptions page of Firstboot when LANG was set to ja_JP. This was because the rhn_review_gui screen and the rshm_confirm_subs screens had the same translated name in Japanese. This bug has been fixed and the Confirm Subscriptions page now correctly displays the Subscription list in the aforementioned case.

BZ#830269

When going back from the Subscription list in Firstboot, instead of being directed to the Entitlement Platform Registration page, the Set Up Software Updates page was displayed when LANG was set to ja_JP. This was because the rhn_review_gui screen and the rshm_confirm_subs screens had the same translated name in Japanese. This bug has been fixed and going back from the Subscription list page now works as expected in the scenario described.

All users of subscription-manager are advised to upgrade to these updated packages, which fix these bugs.

5.314.4. RHBA-2012:1008 – subscription-manager bug fix update

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

The Subscription Manager tool allows users to understand the specific products which have been installed on their machines, and the specific subscriptions which their machines are consuming.

Bug Fix

BZ#829426

Messages provided previously by the yum plug-in were insufficient. The yum plug-in has been improved to display more verbose information and warning messages to users (for example when no subscriptions have been consumed, subscriptions have expired or are not active).

All users of subscription-manager are advised to upgrade to these updated packages, which fix this bug.

5.314.5. RHBA-2012:0804 – subscription-manager bug fix and enhancement update

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

The Subscription Manager tool allows users to understand the specific products which have been installed on their machines, and the specific subscriptions which their machines are consuming.

Enhancements

BZ#768419

The User Interface used generic searching and was not tied any workflow. With this update, the UI has been restructured to focus on service levels. Once a user selects a service level, the rest of the features use that service level to simplify the workflow, making the UI more intuitive for users.

BZ#767620

Subscription Manager previously assumed that users requested content from the Subscription Service only (via the redhat.repo). Users who wished to use their own sources (and repos) had to manually disable the repo provided by Subscription Manager. This update adds a new configuration option which disables the creation of repo files. This option, "manage_repos=0", can be configured in the "/etc/rhsm/rhsm.conf" file. Users are now able to use subscription-manager for subscription management only, and not use the content access features.

BZ#749433

This update provides a new tool to migrate RHN Classic customers to the certificate-based RHN: "rhn-migrate-classic-to-rhsm". For more information on this tool, refer to section "Migrating Systems

from RHN Classic to Certificate-based Red Hat Network" in the Red Hat Enterprise Linux 6 Deployment Guide.

BZ#782433

To register against the customer portal, users had to provide username and password credentials, which could then be logged or kept in the history collection and potentially cause a security issue. With the introduction of Subscription Asset Manager, users can now use activation keys as a replacement for username/password credentials.

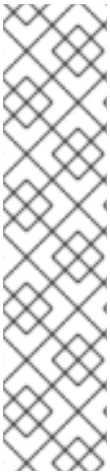
Users are advised to upgrade to these updated subscription-manager packages, which resolve these issues and add these enhancements.

5.314.6. RHBA-2013:1388 – subscription-manager and python-rhsm bug fix and enhancement update

Updated python-rhsm packages and subscription-manager 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-rhsmpackages have been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version. (BZ#967460)
- The subscription-manager packages have been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version. (BZ#967476)
- The subscription-manager-migration-data packages have been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version. (BZ#967478)

Bug Fixes

BZ#967460

The python-rhsm packages have been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version.

BZ#967476

The subscription-manager packages have been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version.

BZ#967478

The subscription-manager-migration-data packages have been upgraded to the latest upstream version, which provides a number of bug fixes and enhancements over the previous version.

The subscription-manager packages provide programs and libraries to allow users to manage subscriptions and yum repositories from the Red Hat Entitlement platform.

Users of subscription-manager and python-rhsm are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.315. SUBVERSION AND NEON

5.315.1. RHEA-2012:0896 — subversion and neon bug fix and enhancement update

Updated subversion and neon packages that fix several bugs and add an enhancement are now available.

Subversion (SVN) is a concurrent version control system which enables one or more users to collaborate in developing and maintaining a hierarchy of files and directories while keeping a history of all changes.

Neon is an HTTP library and WebDAV client library used by Subversion.

Bug Fixes

BZ#749494

The "svn" command unnecessarily required access to the parent directory during certain types of merge operations, which could be denied by the authorization policy on the server. The SVN client has been fixed to not require such access.

BZ#751321

When the "AuthzForceUsernameCase lower" directive was configured in the /etc/httpd/conf.d/subversion.conf file, the "mod_authz_svn" module could crash with a segmentation fault. With this update, segmentation faults no longer occur when using the "AuthzForceUsernameCase" directive.

BZ#798636

Due to a bug in the neon HTTP library, the Server Name Indication (SNI) support was disabled on an SVN client. This update upgrades the neon library, and SNI now works as expected.

Enhancement

BZ#711904, BZ#720790

This update adds an init script for the "svnserve" daemon.

Users are advised to upgrade to these updated subversion and neon packages, which resolve these issues and add this enhancement.

5.316. SUDO

5.316.1. RHSA-2012:1081 — Moderate: sudo security update

An updated sudo package that fixes one security issue is now available for Red Hat Enterprise Linux 5 and 6.

The Red Hat Security Response Team has rated this update as having moderate security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) 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 Fix

CVE-2012-2337

A flaw was found in the way the network matching code in sudo handled multiple IP networks listed in user specification configuration directives. A user, who is authorized to run commands with sudo on specific hosts, could use this flaw to bypass intended restrictions and run those commands on hosts not matched by any of the network specifications.

All users of sudo are advised to upgrade to this updated package, which contains a backported patch to correct this issue.

5.316.2. RHBA-2012:1142 – sudo bug fix update

Updated sudo packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The sudo (superuser do) utility allows system administrators to give certain users the ability to run commands as root.

Bug Fix

BZ#840872

Due to a regression, the suspend and resume actions of commands in sudo shell did not work properly. Consequently, the sudo shell could become unresponsive. A patch has been provided to address this issue and commands in sudo shell can now be suspended and resumed as expected.

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

5.316.3. RHBA-2012:0905 – sudo bug fix update

Updated sudo packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The sudo packages provide the superuser do (sudo) utility, which allows system administrators to give certain users the ability to run commands as root.

Bug Fixes

BZ#604297

Previously, the "-c" check used a very restrictive policy and "visudo -s" treated unused aliases as errors. This update modifies this behavior and "visudo -s" only warns about unused aliases.

BZ#667120

Previously, core dumping in sudo was disabled in the code. Administrators could not control the core dumping. This update modifies the code so that core dumping is no any longer disabled. Now, administrators can control core dumping in sudo, which is a SUID binary, using the /proc/sys/fs/suid_dumpable file.

BZ#697775

Previously, the "sudoedit" used the wrong SELinux context when manipulating files. Files could not be edited when SELinux was in enforcing mode, if the sudoers rule specified a SELinux context that permitted sudoedit. This update modifies the code to permit a transition to the correct SELinux context. Now, files can be edited using the correct SELinux context.

BZ#751680

Previously, the alias checking code in sudo caused false negatives and positives. Syntactically correct sudoers files were declared to be erroneous and unused aliases were not detected. This update modifies the checking code to eliminate false positives and negatives.

BZ#760843

Previously, The nslcd service could not be started if the nscld.conf file contained sudo specific configuration directives. The nslcd daemon could not run while the LDAP sudoers sources were configured. This update uses the separate sudo-ldap config file for configuring LDAP sudoers sources.

BZ#769701

Previously, sudo could handle signals incorrectly if the SIGCHLD signal was received immediately before the select()call and the sudo process became unresponsive after receiving the SIGCHLD signal. This update modifies the underlying code to improve the signal handling.

BZ#797511

Previously, the getgrouplist() function checked the invoker's group membership instead of the membership of the specified user. As a Consequence, sudo listed privileges granted to any group the invoking user was a member of when attempting to view all allowed and forbidden commands both for the invoking user with the "-l" option and for users specified by the "-U" option. This update modifies the getgrouplist() function to correctly check the group membership of the intended user.

BZ#806095

Previously, sudo escaped non-aplhanumeric characters in commands using "sudo -s" or "sudo -" at the wrong place and interfered with the authorization process. Some valid commands were not permitted. Now, non-aplhanumeric characters escape immediately before the command is executed and no longer interfere with the authorization process.

BZ#810147

Previously, the sudo tool interpreted a Runas alias that specified a group incorrectly as a user alias. As a consequencee, the alias appeared to be ignored. This update modifies the code to interprete these aliases and the Runas group aliases are honored as expected.

BZ#810326

Previously, the sudo word wrapping feature caused output to be wrapped at terminal width boundary even in output that was piped to an other command. This update modifies the underlying code to detect whether the output is a pipe and disables the word wrapping feature in this case.

BZ#810372

Previously, the "tls_checkpeer" option was set on a handle that is not used when connecting to the Lightweight Directory Access Protocol (LDAP) server. The "tls_checkpeer" option could not be disabled. This update modifies the underlying code so that the option can now be disabled.

All users of sudo are advised to upgrade to this updated package, which fix these bugs.

5.316.4. RHBA-2013:0619 – sudo bug fix update

Updated sudo packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The superuser do (Sudo) utility allows system administrators to give certain users the ability to run commands as root.

Bug Fix

BZ#891593

Previously, the sudo utility executed commands directly and replaced the sudo process. However, in a previous update, internal execution method of commands in sudo changed and sudo now runs commands as child processes. This change in behavior caused problems with custom scripts. This update adds the cmnd_no_wait option; with it, the old behavior is restored and commands are executed directly in the sudo process, thus fixing this bug.

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

5.317. SYSFSUTILS

5.317.1. RHBA-2012:1453 – sysfsutils bug fix update

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

The sysfsutils packages provide a suite of daemons to manage access to remote directories and authentication mechanisms. The sysfsutils suite provides an NSS and PAM interface toward the system and a pluggable backend system to connect to multiple different account sources. It is also the basis to provide client auditing and policy services for projects like FreeIPA.

Bug Fix

BZ#671554

Prior to this update, sysfs directories were not closed as expected. As a consequence, the libsysfs library could leak memory in long running programs that frequently opened and closed sysfs directories. This update modifies the underlying code to close sysfs directories as expected.

All users of sysfsutils are advised to upgrade to these updated packages, which fix this bug.

5.318. SYSLINUX

5.318.1. RHBA-2012:0894 – syslinux bug fix update

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

The syslinux utility is responsible for booting the operating system kernel.

Bug Fixes

BZ#729013

Prior to this update, the binaries were compiled without debugging information. As a consequence, the debuginfo files for some binaries were missing. With this update, the binaries include all required debug information.

BZ#806464

Prior to this update, the isohybrid utility failed with a seek error when attempting to write ISO images to a USB device. As a consequence, the ISO images were on some systems not bootable. This update modifies the underlying code so that isohybrid now successfully writes the images to USB devices.

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

5.319. SYSSTAT

5.319.1. RHBA-2012:0866 – sysstat bug fix and enhancement update

An updated sysstat package that fixes multiple bugs and adds three enhancements is now available for Red Hat Enterprise Linux 6.

The sysstat package contains a set of utilities which enable system monitoring of disks, network, and other I/O activity.

Bug Fixes

BZ#690562

Prior to this update, the cifsiostat utility did not report the correct number of open files on CIFS file systems. A patch has been applied to ensure that the cifsiostat utility outputs the correct number of open files on CIFS file systems.

BZ#694759

Prior to this update, the "iostat -n" command did not list NFS shares with excessively long names. The underlying source code has been revised to ensure that the command lists all NFS shares available.

BZ#771594

Previously, nr_requests could overflow if set to high values. As a consequence, the "iostat" and "sar -d" command would report the overflowed values from the /proc/diskstats file instead of the proper values. The underlying source code has been modified to ensure that the "iostat" and "sar -d" commands output the correct information.

BZ#801453

Previously, the sa2 script could return an error code even when the script completed successfully. The underlying source code has been modified so that the sa2 script now returns the correct code on successful completion.

BZ#801702

Running the "sar" command with the "-p" option to "pretty-print" device names did not work as expected if the device minor number was greater than 256. The device names were incorrectly displayed as "nodev" or "dev-[major number]-[minor number]". With this update, the value of the

IOC_MAXNIMOR constant has been increased, and devices with the minor number greater than 256 are now displayed correctly in the "sar" output.

Enhancements

BZ#674648

Previously, on 64-bit PowerPC architectures, the mpstat utility displayed all processors, including both busy and idle processors with zero activity. This update introduces a new option for the mpstat utility, "-P ON". If this option is used, mpstat lists only online processors.

BZ#693398

With this update, the SADC_OPTIONS configuration variable has been moved from the sysstat init script, located in the /etc/init.d/ directory, to the sysstat configuration file located in the /etc/sysconfig/ directory. Also, a note about the -S option was added to the sadc manual page.

BZ#766431

Previously, the iostat utility did not display target device information if a symbolic link was specified as an input parameter. This update adds support for symbolic links as input parameters in the iostat utility.

All users of sysstat are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.320. SYSTEM-CONFIG-DATE-DOCS

5.320.1. RHBA-2012:0934 – system-config-date-docs bug fix update

Updated system-config-date-docs packages that fix one bug are now available for Red Hat Enterprise Linux 6.

The system-config-date-docs packages contain the online documentation for system-config-date, with which you can configure date, time and the use of time servers on your system.

Bug Fix

BZ#691572

Prior to this update, the help documentation contained out-of-date screenshots and the text did not correctly reflect the user interface elements. This update contains updated screenshots and documents the user interface correctly.

All users are advised to upgrade to these updated packages, which fix this bug.

5.321. SYSTEM-CONFIG-KDUMP

5.321.1. RHBA-2012:0824 – system-config-kdump bug fix and enhancement update

Updated system-config-kdump packages that fix multiple bugs and add three enhancements are now available for Red Hat Enterprise Linux 6.

The system-config-kdump packages provide a graphical tool to configure kernel crash dumping via kdump and kexec.

Bug Fixes

BZ#590057

Previously, the system-config-kdump tool did not handle the kexec service status correctly. As a consequence, the system-config-kdump service could fail at start. This update synchronizes the system-config-kdump and kdump service activation. Now, the service is started, stopped or restarted as expected.

BZ#609487

Previously, the system-config-kdump tool used synchronous dbus commands. This update uses asynchronous dbus commands. Now, the system-config-kdump tool waits for the processes running in the background.

BZ#626787

Previously, system-config-kdump used a file-chooser button as well as input via a text box to specify a path for kdump, which could cause confusion. This update uses only text box input and also informs the user of the location of the copied kdump.

BZ#629483

Previously, The "About" dialog used an incorrect version number. This update modifies the dialog so that the right version number is now displayed.

BZ#632999

Previously, not all strings in the POT file were marked for localization. As a consequence, several dialogs could not be translated. This update marks the missing strings in source files to make the POT file complete. .

BZ#642751

Previously, a dialog box with the error message "Core collector must begin with makedumpfile multiplied and could not be closed" if the kdump.conf file contained the "core_collector cp" command. This update modifies the underlying code so that the dialog box is shown only once.

BZ#653450

Previously, several system-config-kdump messages contained misprints. This update modifies the strings and all messages are now correct.

BZ#676777

Previously, the system-config-kdump tool was not constructed to accept more than one value. As a consequence, the extended crashkernel syntax was not handled correctly. This update modifies the underlying code so that system-config-kdump can now read the extended syntax, but always writes in basic syntax.

BZ#740155

Previously, the system-config-kdump tool wrongly showed an error message when values from the /proc/iomem file could not be read on 64-bit PowerPC platforms. With this update, only an informational message is shown.

BZ#754059

Previously, the system-config-kdump tool used the wrong format to save the target type nfs. With this update, the nfs network target is saved correctly.

BZ#813337

Previously, the system-config-kdump tool failed to configure zipl on IBM S/390 systems.This update modifies the zipl helper script so that all configurations are now correctly updated.

BZ#821410

Previously, the error message "module" object has no attribute "show_call_call_error_message" contained a misprint. This update removes the second "call".

BZ#819814

Previously, the system-config-kdump utility contained various locales that were not completely translated. This update adds the missing translations for supported languages.

Enhancements

BZ#622870

Previously, the system-config-kdump tool did not display messages about dbus errors. This update adds meaningful dbus error messages.

BZ#796308

Previously, the system-config-kdump tool did not support IBM S/390 hardware. This update modifies the installer so that IBM S/390 hardware is now enabled.

BZ#816009

The threshold for allowing auto configuration of kernel dumping was changed to 2 GB. With this update, system-config-kdump reflects this change.

All users of system-config-kdump are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.322. SYSTEM-CONFIG-KEYBOARD

5.322.1. RHEA-2012:0852 – system-config-keyboard enhancement update

Updated system-config-keyboard packages that add one enhancement 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.

Enhancement

BZ#771389

Prior to this update, the Red Hat Enterprise Virtualization Hypervisor pulled too many dependencies from the system-config-keyboard package to support keyboard selection capability for non-US keyboards. This update adds the system-config-keyboard-base package that contains the core python libraries.

All users of system-config-keyboard are advised to upgrade to these updated packages, which add this enhancement.

5.323. SYSTEM-CONFIG-LANGUAGE

5.323.1. RHBA-2012:1213 – system-config-language bug fix update

An updated system-config-language package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The system-config-language is a graphical user interface that allows the user to change the default language of the system.

Bug Fix

BZ#819811

When using system-config-language in a non-English locale, some of the messages in the GUI were not translated. Consequently, non-English users were presented with untranslated messages. With this update, all message strings have been translated.

All users of system-config-language are advised to upgrade to this updated package, which fixes this bug.

5.324. SYSTEM-CONFIG-LVM

5.324.1. RHBA-2012:0960 – system-config-lvm bug fix update

An updated system-config-lvm package that fixes multiple bugs is now available for Red Hat Enterprise Linux 6.

The system-config-lvm package contains a utility for configuring logical volumes using a graphical user interface.

Bug Fixes

BZ#791153

The system-config-lvm utility incorrectly displayed all LVM devices in the "Unitialized Entities" group. This could confuse users who could try to re-initialize logical volumes unnecessarily. This update modifies system-config-lvm to show LVM devices as "Unitialized Entities" only if the logical volumes really are uninitialized.

BZ#708029

The system-config-lvm utility uses fdisk as a back end for partition operations. The fdisk utility does not support the Extensible Firmware Interface (EFI) GUID Partition Table (GPT), so an attempt to initialize a physical partition on EFI GPT fails. However, system-config-lvm previously did not handle this situation correctly and the operation failed without any error message. This update modifies the underlying code so that system-config-lvm now provides a valid error message in this situation.

BZ#726830

The system-config-lvm utility did not properly handle physical partitions with names in the form of "*p[0-9]", for example, "loop0p0". Therefore, system-config-lvm terminated unexpectedly when

attempting to initialize such a partition. This update modifies the underlying code so that system-config-lvm no longer crashes in this scenario. However, system-config-lvm still does not handle such partition names and such partitions cannot be initialized.

BZ#815921

The system-config-lvm utility uses the dmsetup tool for obtaining information from Device Mapper Multipath. Output of dmsetup was designed to be parsed by scripts, however, with an update of device-mapper-multipath to version 1.02.74, this output was changed. As a consequence, system-config-lvm terminated unexpectedly with a traceback when used with Device-Mapper Multipath available on the system. This update modifies system-config-lvm to use a user-specified output which does not change across different versions. The system-config-lvm utility now works as expected.

All users of system-config-lvm are advised to upgrade to this updated package, which fixes these bugs.

5.325. SYSTEM-CONFIG-PRINTER

5.325.1. RHBA-2012:0448 – system-config-printer bug fix update

Updated system-config-printer packages that fix two bugs are now available for Red Hat Enterprise Linux 6.

The system-config-printer package contains a print queue configuration tool with a graphical user interface.

Bug Fixes

BZ#739745

Previously, displaying tooltips while being in a main loop recursion caused the system-config-printer utility to terminate unexpectedly. To prevent system-config-printer from crashing, displaying tooltips is now avoided during the main loop recursion.

BZ#744519

Python bindings for the CUPS library were not reliable when threads were used. In particular, a single password callback function was used instead of one for each thread. This, in some cases, caused the system-config-printer utility to terminate unexpectedly with a segmentation fault. With this update, thread local storage is used for the password callback function in Python bindings for the CUPS library.

All users of system-config-printer are advised to upgrade to these updated packages, which fix these bugs.

5.326. SYSTEM-CONFIG-USERS

5.326.1. RHBA-2012:1387 – system-config-users bug fix update

Updated system-config-users packages that fix three bugs are now available for Red Hat Enterprise Linux 6.

The system-config-users packages provide a graphical utility for administrating users and groups.

Bug Fixes

BZ#736037

Prior to this update, expiration dates at or before January 1, 1970 were not correctly calculated. As a consequence, the system-config-users utility stored expiration dates off by one day into /etc/shadow. This update modifies the underlying code so that account expiration dates are calculated and stored correctly.

BZ#801652

Prior to this update, a string in the user interface was not correctly localized into Japanese. This update modifies the string so that the text is now correct.

BZ#841886

Prior to this update, the system-config-users utility determined incorrectly whether to set an account as inactive if an expired password was not reset during a specified period. This update modifies the underlying code to check for this condition by hard-coding the value which indicates this condition.

All users of system-config-users are advised to upgrade to these updated packages, which fix these bugs.

5.327. SYSTEMTAP

5.327.1. RHBA-2012:1337 – systemtap bug fix update

Updated systemtap packages that fix a bug 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.

Bug Fix

BZ#859832

In previous kernels and versions of systemtap, the nfsd.open probe-alias in the nfsd tapset referred to the "access" parameter, which was later renamed to "may_flags" in the kernel. Consequently, the semantic errors occurred and then the stap command failed to execute. This update lets the nfsd.open probe-alias check under both names for setting the "access" script-level variable, and stap now works as expected in the described scenario.

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

5.327.2. RHBA-2012:0820 – systemtap bug fix and enhancement update

Updated systemtap packages that fix multiple bugs and add several 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.

BZ#751478

The systemtap packages have been upgraded to upstream version 1.7, which provides a number of bug fixes and enhancements.

Bug Fixes

BZ#588359

SystemTap did not fully implement backtraces for the PowerPC and the IBM System z platforms and the backtraces were not always informative. Now, the support for PowerPC and IBM System z backtraces is improved.

BZ#639338

The custom functions "deref()" and "store_deref()" in the SystemTap support code did not correctly handle the memory access for the IBM System z platform and some SystemTap tests could fail. Now, the SystemTap runtime use the kernel version of these functions.

BZ#738365

The usymbols.exp test did not correctly cast some arguments for data structures it accessed and subsequently failed on PowerPC and IBM System z platforms. This update modifies the test to successfully run on PowerPC and IBM System z platforms.

BZ#752170

The SystemTap translator did not parse code with user-space markers containing some types of operand location information. Now, the SystemTap translator understands the additional x86_64 address modes.

BZ#752568

Compile servers took a long time to translate scripts into instrumentation and send them back. Now, the verbose options provide information about the progress of the compiler server to verify that the compiler server is working as expected.

BZ#754567

The systemtap-client script functionality was folded into the systemtap package, eliminating the systemtap-client package. As a consequence, a client-only systemtap installation was not available. This update provides the systemtap-client package to allow for client-only systemtap installations.

BZ#790091

The stap-serverd daemon hard-coded some rlimit values when running under the stap-server User ID. The daemon failed if the rlimit values were exceeded. Now, stap-serverd refers to configuration information stored in ~/stap-server/.systemtap/rc. .

BZ#791243

The tcp.sendmsg probe alias variable sock was not correctly set for certain kernel versions. Scripts that used the sock variable when probing tcp.sendmsg did not compile. Now, the tcp.sendmsg alias initialization of the sock variable handles these variations in the kernel structure.

BZ#812871

The nd_syscall tests used several probes that applied a sixth argument. Consequently, these probes could fail on IBM System z platforms which did not support access to a sixth argument if a function was on the stack. Now, the access code for the IBM System z runtime is improved to allow access to arguments on the stack.

BZ#813323

SystemTap did not run precompiled scripts if users with membership only of the stapusr group attempted to run an unsigned precompiled script from the /lib/modules/`uname -r`/systemtap directory. Now, SystemTap elevates the privileges for precompiled scripts in the systemtap directory.

Enhancement

BZ#798754

This update adds the enospc.stp script to obtain a direct information when a file system is full.

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

5.328. TAR

5.328.1. RHBA-2012:1372 – tar bug fix update

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

The tar packages provide the GNU tar program. Gnu tar can allows to save multiple files in one archive and can restore the files from that archive. This update fixes the following bug:

Bug Fix

BZ#841308

Prior to this update, tar failed to match and extract given file names from an archive when this archive was created with the options "--sparse" and "--posix". This update modifies the underlying code to match and extract the given name as expected.

All users of tar are advised to upgrade to these updated packages, which fix this bug.

5.328.2. RHBA-2012:0849 – tar bug fix and enhancement update

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

The GNU tar program can save multiple files in one archive and restore the files from that archive.

Bug Fixes

BZ#653433

Before this update, tar could terminate with a segmentation fault and returned code 139. This happened when tar was used for incremental backup of the root directory with the option "--listed-incremental" (short option "-g") due to incorrect directory name resolution. The root directory name is now resolved correctly and the backup process succeeds in this scenario.

BZ#656834

The tar utility archived sparse files with long names (about 100 characters) incorrectly if run with the "--posix" and "--sparse" options (PAX mode). Such files were stored with misleading names inside the tar archive as there was not enough space allocated for the file names. Subsequent unpacking of

the package resulted in confusing output file names. With this update, more space is now allocated for the file names in this scenario and the problem no longer occurs.

BZ#698212

If tar was run with the "--remove-files" option and the archived directory contained a file and a symbolic link pointing to the file, the file was deleted but not backed up. The archiving process terminated with an error. With this update, the file is archived as expected in this scenario.

BZ#726723

The tar unpacking process could enter an infinite loop and consume extensive CPU resources when run with the "--keep-old-files" option. This happened when unpacking an archive with symbolic links and the target of the symbolic link already existed. With this update, the code has been modified to handle symbolic links correctly in this scenario.

BZ#768724

The tar tool used the glibc fnmatch() function to match file names. However, the function failed to match a file name when the archived file name contained characters not supported by the default locale. Consequently, the file was not unpacked. With this update, tar uses the gnulib fnmatch() and the file name is matched as expected.

BZ#782628

If tar was run with the "--remove-files" option, it failed to remove the archived files when append mode was activated (the -r option). With this update, tar with the "--remove-files" option now calls the function that removes the files after they have been archived and the option works as expected.

BZ#688567

The tar tool failed to update the target archive when run with the "--update" and "--directory" options, returned the "Cannot stat: No such file or directory" error message, and the directory content was not archived. With this update, the tar command with the two options now works as expected.

BZ#799252

When extracting an archive with the "--keep-old-files" option, tar silently skipped already existing files. With this update, tar returns error code 2 and a warning in this scenario. Also, the "--skip-old-files" option has been added to allow the previous "--keep-old-files" behavior without returning errors for files that already exist.

BZ#807728

When run with the "--list" (-r) option, tar returned the "tar: write error" message, even though the execution succeeded. This happened if the command used redirection with a pipeline and the command following the redirection failed to process the entire tar command output. With this update, the spurious message is no longer returned in this scenario.

Enhancement

BZ#760665

When archiving a sparse file containing 0 blocks of data, the archiving process experienced severe performance issues because tar was scanning the sparse file for non-existing data. With this update, a sparse file containing 0 blocks is detected by the stat() call and the archiving process is now faster for such files.

All tar users are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.329. TBOOT

5.329.1. RHBA-2012:0771 – tboot bug fix update

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

The tboot packages provide Trusted Boot (tboot), an open source pre- kernel/VMM module, that uses Intel Trusted Execution Technology (Intel TXT) to initialize the launch of operating system kernels and virtual machines.

The tboot package has been upgraded to upstream version 1.7.0, which provides a number of bug fixes and enhancements over the previous version and adds support for future Intel chip sets. (BZ#773406)

Bug Fixes

BZ#732439

Prior to this update, increasing binary large objects (BLOB) with lcptools or lcputils utilities could cause a null pointer dereferencing. This update modifies the underlying code so that the pointer now refers to an existing value.

BZ#754345

Prior to this update, the tboot module could, under certain circumstances, be enabled on 32-bit platforms. As a consequence, tboot could prevent the kernel from booting. This update restricts tboot to 64-bit platforms.

All users of tboot are advised to upgrade to these updated packages, which fix these bugs and adds these enhancements.

5.330. TCPDUMP

5.330.1. RHBA-2012:0414 – tcpdump bug fix update

An updated tcpdump package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The tcpdump tool is a command line utility for monitoring network traffic.

Bug Fix

BZ#684005

Previously, the "start-time" command-line argument of the tcpslice utility was parsed incorrectly. As a consequence, the utility produced error messages every time the command-line argument was used. With this update, the "start-time" command-line argument is parsed correctly, and an error message is displayed only if "start-time" is defined in an incorrect format.

All users of tcpdump are advised to upgrade to this updated package, which fixes this bug.

5.331. TELNET

5.331.1. RHBA-2012:1312 – telnet bug fix update

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

Telnet is a popular protocol for logging in to remote systems over the Internet. The telnet-server package includes a telnet service that supports remote logins into the host machine. The telnet service is disabled by default.

Bug Fix

BZ#860012

Prior to this update, the telnetd daemon did not reuse previously created entries in the /var/run/utmp file. Consequently, /var/run/utmp grew and contained empty entries, eventually causing various other problems. With this update, telnetd has been fixed to reuse entries in /var/run/utmp and behave like other programs which use this file, thus preventing this bug.

All users of telnet are advised to upgrade to these updated packages, which fix this bug.

5.332. THUNDERBIRD

5.332.1. RHSA-2012:1351 – Critical: thunderbird security update

An updated thunderbird package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fixes

CVE-2012-3982, CVE-2012-3988, CVE-2012-3990, CVE-2012-3995, CVE-2012-4179, CVE-2012-4180, CVE-2012-4181, CVE-2012-4182, CVE-2012-4183, CVE-2012-4185, CVE-2012-4186, CVE-2012-4187, CVE-2012-4188

Several flaws were found in the processing of malformed content. Malicious content could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-3986, CVE-2012-3991

Two flaws in Thunderbird could allow malicious content to bypass intended restrictions, possibly leading to information disclosure, or Thunderbird executing arbitrary code. Note that the information disclosure issue could possibly be combined with other flaws to achieve arbitrary code execution.

CVE-2012-1956, CVE-2012-3992, CVE-2012-3994

Multiple flaws were found in the location object implementation in Thunderbird. Malicious content could be used to perform cross-site scripting attacks, script injection, or spoofing attacks.

CVE-2012-3993, CVE-2012-4184

Two flaws were found in the way Chrome Object Wrappers were implemented. Malicious content could be used to perform cross-site scripting attacks or cause Thunderbird to execute arbitrary code.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Christian Holler, Jesse Ruderman, Soroush Dalili, miaubiz, Abhishek Arya, Atte Kettunen, Johnny Stenback, Alice White, moz_bug_r_a4, and Mariusz Mlynski as the original reporters of these issues.

Note: None of the issues in this advisory can be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 10.0.8 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.2. RHSA-2012:1362 – Critical: thunderbird security update

An updated thunderbird package that fixes one security issue is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fix

CVE-2012-4193

A flaw was found in the way Thunderbird handled security wrappers. Malicious content could cause Thunderbird to execute arbitrary code with the privileges of the user running Thunderbird.

Red Hat would like to thank the Mozilla project for reporting this issue. Upstream acknowledges moz_bug_r_a4 as the original reporter.



NOTE

This issue cannot be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. It could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which corrects this issue. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.3. RHSA-2012:1413 – Important: thunderbird security update

An updated thunderbird package that fixes multiple security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fix

CVE-2012-4194, CVE-2012-4195, CVE-2012-4196

Multiple flaws were found in the location object implementation in Thunderbird. Malicious content could be used to perform cross-site scripting attacks, bypass the same-origin policy, or cause Thunderbird to execute arbitrary code.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Mariusz Mlynski, moz_bug_r_a4, and Antoine Delignat-Lavaud as the original reporters of these issues.



NOTE

None of the issues in this advisory can be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 10.0.10 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.4. RHSA-2012:1089 – Critical: thunderbird security update

An updated thunderbird package that fixes multiple security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fixes

CVE-2012-1948, CVE-2012-1951, CVE-2012-1952, CVE-2012-1953, CVE-2012-1954, CVE-2012-1958, CVE-2012-1962, CVE-2012-1967

Several flaws were found in the processing of malformed content. Malicious content could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-1959

Malicious content could bypass same-compartment security wrappers (SCSW) and execute arbitrary code with chrome privileges.

CVE-2012-1955

A flaw in the way Thunderbird called history.forward and history.back could allow an attacker to conceal a malicious URL, possibly tricking a user into believing they are viewing trusted content.

CVE-2012-1957

A flaw in a parser utility class used by Thunderbird to parse feeds (such as RSS) could allow an attacker to execute arbitrary JavaScript with the privileges of the user running Thunderbird. This issue could have affected other Thunderbird components or add-ons that assume the class returns sanitized input.

CVE-2012-1961

A flaw in the way Thunderbird handled X-Frame-Options headers could allow malicious content to perform a clickjacking attack.

CVE-2012-1963

A flaw in the way Content Security Policy (CSP) reports were generated by Thunderbird could allow malicious content to steal a victim's OAuth 2.0 access tokens and OpenID credentials.

CVE-2012-1964

A flaw in the way Thunderbird handled certificate warnings could allow a man-in-the-middle attacker to create a crafted warning, possibly tricking a user into accepting an arbitrary certificate as trusted.

The nss update RHBA-2012:0337 for Red Hat Enterprise Linux 5 and 6 introduced a mitigation for the CVE-2011-3389 flaw. For compatibility reasons, it remains disabled by default in the nss packages. This update makes Thunderbird enable the mitigation by default. It can be disabled by setting the NSS_SSL_CBC_RANDOM_IV environment variable to 0 before launching Thunderbird. (BZ#838879)

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Benoit Jacob, Jesse Ruderman, Christian Holler, Bill McCloskey, Abhishek Arya, Arthur Gerkis, Bill Keese, moz_bug_r_a4, Bobby Holley, Mariusz Mlynski, Mario Heiderich, Frédéric Buclin, Karthikeyan Bhargavan, and Matt McCutchen as the original reporters of these issues.

Note: None of the issues in this advisory can be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 10.0.6 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.5. RHSA-2013:0145 – Critical: thunderbird security update

An updated thunderbird package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fixes

CVE-2013-0744, CVE-2013-0746, CVE-2013-0750, CVE-2013-0753, CVE-2013-0754, CVE-2013-0762, CVE-2013-0766, CVE-2013-0767, CVE-2013-0769

Several flaws were found in the processing of malformed content. Malicious content could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2013-0758

A flaw was found in the way Chrome Object Wrappers were implemented. Malicious content could be used to cause Thunderbird to execute arbitrary code via plug-ins installed in Thunderbird.

CVE-2013-0759

A flaw in the way Thunderbird displayed URL values could allow malicious content or a user to perform a phishing attack.

CVE-2013-0748

An information disclosure flaw was found in the way certain JavaScript functions were implemented in Thunderbird. An attacker could use this flaw to bypass Address Space Layout Randomization (ASLR) and other security restrictions.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Atte Kettunen, Boris Zbarsky, pa_kt, regenrecht, Abhishek Arya, Christoph Diehl, Christian Holler, Mats Palmgren, Chiaki Ishikawa, Mariusz Mlynski, Masato Kinugawa, and Jesse Ruderman as the original reporters of these issues.



NOTE

All issues except CVE-2013-0744, CVE-2013-0753, and CVE-2013-0754 cannot be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 10.0.12 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.6. RHSA-2012:1483 – Critical: thunderbird security update

An updated thunderbird package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fixes

CVE-2012-4214, CVE-2012-4215, CVE-2012-4216, CVE-2012-5829, CVE-2012-5830, CVE-2012-5833, CVE-2012-5835, CVE-2012-5839, CVE-2012-5840, CVE-2012-5842

Several flaws were found in the processing of malformed content. Malicious content could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-4202

A buffer overflow flaw was found in the way Thunderbird handled GIF (Graphics Interchange Format) images. Content containing a malicious GIF image could cause Thunderbird to crash or, possibly, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-4207

A flaw was found in the way Thunderbird decoded the HZ-GB-2312 character encoding. Malicious content could cause Thunderbird to run JavaScript code with the permissions of different content.

CVE-2012-4209

A flaw was found in the location object implementation in Thunderbird. Malicious content could possibly use this flaw to allow restricted content to be loaded by plug-ins.

CVE-2012-5841

A flaw was found in the way cross-origin wrappers were implemented. Malicious content could use this flaw to perform cross-site scripting attacks.

CVE-2012-4201

A flaw was found in the evalInSandbox implementation in Thunderbird. Malicious content could use this flaw to perform cross-site scripting attacks.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Abhishek Arya, miaubiz, Jesse Ruderman, Andrew McCreight, Bob Clary, Kyle Huey, Atte Kettunen, Masato Kinugawa, Mariusz Mlynski, Bobby Holley, and moz_bug_r_a4 as the original reporters of these issues.



NOTE

All issues except CVE-2012-4202 cannot be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 10.0.11 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.7. RHSA-2012:1211 – Critical: thunderbird security update

An updated thunderbird package that fixes multiple security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fixes

CVE-2012-1970, CVE-2012-1972, CVE-2012-1973, CVE-2012-1974, CVE-2012-1975, CVE-2012-1976, CVE-2012-3956, CVE-2012-3957, CVE-2012-3958, CVE-2012-3959, CVE-2012-3960, CVE-2012-3961, CVE-2012-3962, CVE-2012-3963, CVE-2012-3964

Several flaws were found in the processing of malformed content. Malicious content could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-3969, CVE-2012-3970

Content containing a malicious Scalable Vector Graphics (SVG) image file could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-3967, CVE-2012-3968

Two flaws were found in the way Thunderbird rendered certain images using WebGL. Malicious content could cause Thunderbird to crash or, under certain conditions, possibly execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-3966

A flaw was found in the way Thunderbird decoded embedded bitmap images in Icon Format (ICO) files. Content containing a malicious ICO file could cause Thunderbird to crash or, under certain conditions, possibly execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-3980

A flaw was found in the way the "eval" command was handled by the Thunderbird Error Console. Running "eval" in the Error Console while viewing malicious content could possibly cause Thunderbird to execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2012-3972

An out-of-bounds memory read flaw was found in the way Thunderbird used the format-number feature of XSLT (Extensible Stylesheet Language Transformations). Malicious content could possibly cause an information leak, or cause Thunderbird to crash.

CVE-2012-3978

A flaw was found in the location object implementation in Thunderbird. Malicious content could use this flaw to possibly allow restricted content to be loaded.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Gary Kwong, Christian Holler, Jesse Ruderman, John Schoenick, Vladimir Vukicevic, Daniel Holbert, Abhishek Arya, Frédéric Hoguin, miaubiz, Arthur Gerkis, Nicolas Grégoire, moz_bug_r_a4, and Colby Russell as the original reporters of these issues.

Note: All issues except CVE-2012-3969 and CVE-2012-3970 cannot be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 10.0.7 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.332.8. RHSA-2013:0272 – Critical: thunderbird security update

An updated thunderbird package that fixes several security issues is now available for Red Hat Enterprise Linux 5 and 6.

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

Mozilla Thunderbird is a standalone mail and newsgroup client.

Security Fixes

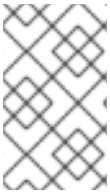
CVE-2013-0775, CVE-2013-0780, CVE-2013-0782, CVE-2013-0783

Several flaws were found in the processing of malformed content. Malicious content could cause Thunderbird to crash or, potentially, execute arbitrary code with the privileges of the user running Thunderbird.

CVE-2013-0776

It was found that, after canceling a proxy server's authentication prompt, the address bar continued to show the requested site's address. An attacker could use this flaw to conduct phishing attacks by tricking a user into believing they are viewing trusted content.

Red Hat would like to thank the Mozilla project for reporting these issues. Upstream acknowledges Nils, Abhishek Arya, Olli Pettay, Christoph Diehl, Gary Kwong, Jesse Ruderman, Andrew McCreight, Joe Drew, Wayne Mery, and Michal Zalewski as the original reporters of these issues.



NOTE

All issues cannot be exploited by a specially-crafted HTML mail message as JavaScript is disabled by default for mail messages. They could be exploited another way in Thunderbird, for example, when viewing the full remote content of an RSS feed.



IMPORTANT

This erratum upgrades Thunderbird to version 17.0.3 ESR. Thunderbird 17 is not completely backwards-compatible with all Mozilla add-ons and Thunderbird plug-ins that worked with Thunderbird 10.0. Thunderbird 17 checks compatibility on first-launch, and, depending on the individual configuration and the installed add-ons and plug-ins, may disable said Add-ons and plug-ins, or attempt to check for updates and upgrade them. Add-ons and plug-ins may have to be manually updated.

All Thunderbird users should upgrade to this updated package, which contains Thunderbird version 17.0.3 ESR, which corrects these issues. After installing the update, Thunderbird must be restarted for the changes to take effect.

5.333. TOG-PEGASUS

5.333.1. RHBA-2012:0953 – tog-pegasus bug fix update

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

The tog-pegasus packages provide OpenPegasus Web-Based Enterprise Management (WBEM) services for Linux. WBEM enables management solutions that deliver increased control of enterprise

resources. WBEM is a platform and resource independent Distributed Management Task Force (DMTF) standard that defines a common information model (CIM) and communication protocol for monitoring and controlling resources from diverse sources.

Bug Fixes

BZ#796191

Previously, with the Single Chunk Memory Objects (SCMO) implementation, empty string values in embedded instances were converted to null values during the embedded CIMInstance to SCMOInstance conversion. This was due to the usage of the _setString() function that set the string size to 0 if the string was empty. This broke functionality of the existing providers. A backported upstream patch uses the _SetBinary() function instead which is already used while setting the string values on the normal SCMOInstance.

BZ#799040

Previously, the tog-pegasus packages did not provide a generic "cim-server", which could be required by packages that do not need a specific implementation of the CIM server as a dependency. With this update, the tog-pegasus packages provide a generic "cim-server" that can be required by such packages.

All users of tog-pegasus are advised to upgrade to these updated packages, which fix these bugs.

5.334. TOMCAT6

5.334.1. RHBA-2013:0137 – tomcat6 bug fix update

Updated tomcat6 packages that fix a bug 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 Fix

BZ#852868

When a web application used its own class loader when compiling JSP, a deadlock in Tomcat WebappClassLoader could occur. This update fixes the synchronization bug and external class loaders no longer interfere with WebappClassLoader.

Users of tomcat6 are advised to upgrade to these updated packages, which fix this bug.

5.334.2. RHBA-2012:0945 – tomcat6 bug fix and enhancement update

Updated tomcat6 packages that fix several bugs and provide an enhancement 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#697968

Previously, in certain cases, if "LANG=fr_FR" or "LANG=fr_FR.UTF-8" was set as an environment variable or in "/etc/sysconfig/tomcat6" on 64-bit PowerPC systems, Tomcat may have failed to start correctly. With this update, Tomcat works as expected when LANG is set to "fr_FR" or "fr_FR.UTF-8".

BZ#701759

The "/usr/sbin/tomcat6" wrapper script used a hard-coded path to the "catalina.out" file, which could have caused problems (such as logging init script output) if Tomcat was being run with a user other than "tomcat" and with CATALINA_BASE set to a directory other than the default. With this update, the wrapper script redirects output to ${CATALINA_BASE}/logs/catalina.out for all "start", "start-security", and "stop" actions.

BZ#748813

Using the URL class coupled with the setChunkedStreamingMode() function caused a null pointer exception error and HTTP response status code 405 was returned. A patch has been applied which adds a check for form data before processing. If the requested body length is zero, a null is returned without further processing. As a result, the error no longer occurs in the scenario described.

BZ#783567

Due to a regression, when a JavaServer Pages (JSP) tag that does not allow JSP Expression Language (EL) expression values (such as struts 2 tags) was used, and one of the attributes was passed a certain value (such as a backslash), the parser threw the following exception:

> According to TLD or attribute directive in tag file, attribute value does not accept any expressions

JSP parsing utilizes the directive attribute "deferredSyntaxAllowedAsLiteral" which determines if deferred statements are treated as literals. The default is false. If true, the "#" sign will not be treated as an escape. This update applies an upstream patch and the problem no longer occurs.

Enhancement

BZ#782400

With this update, the tomcat6 dependency on redhat-lsb has been removed. Red Hat Enterprise Linux tomcat6 strives to have Linux Standards Base (LSB) compliant systemv init scripts. However, Java has been absent from the list of compliant binaries since 2011. Since Tomcat runs in the Java Virtual Machine (JVM), there is little that can be done in addition to the init script compliance. The redhat-lsb dependency can be removed with very little risk.

Users are advised to upgrade to these updated tomcat6 packages, which provide numerous bug fixes and enhancement.

5.335. TRACE-CMD

5.335.1. RHEA-2012:0976 – trace-cmd enhancement update

Updated trace-cmd packages that add one enhancement are now available for Red Hat Enterprise Linux 6.

The trace-cmd packages contain a command line tool that interfaces with ftrace in the kernel.

Enhancement

BZ#632061

This update adds support for the "-i" option that can be used to ignore events. By default, if an event is listed but cannot be found by the trace-cmd utility on the system, the utility exits. This option allows trace-cmd execution to continue even when an event is listed on the command line but cannot

be found on the system.

All users of trace-cmd are advised to upgrade to these updated packages, which add this enhancement.

5.336. TSCLIENT

5.336.1. RHBA-2012:0382 – tsclient bug fix update

An updated tsclient package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The tsclient utility is a GTK2 front end that makes it easy to use the Remote Desktop Protocol client (rdesktop) and vncviewer utilities.

Bug Fix

BZ#734826

When opening an X Display Manager Control Protocol (XDMCP) connection using the tsclient utility, tsclient could terminate unexpectedly with a segmentation fault. A patch has been applied to address this issue, so that an XDMCP connection is now started correctly for the configured host.

All users of tsclient are advised to upgrade to this updated package, which fixes this bug.

5.337. TUNED

5.337.1. RHBA-2012:0924 – tuned bug fix and enhancement update

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

The tuned package contains a daemon that tunes system settings dynamically. It does so by monitoring the usage of several system components periodically.

Bug Fixes

BZ#747210

When the diskdevstat or netdevstat tool was run with wrong command-line arguments, the tool returned a compilation error and exited. Both tools have been fixed to check the command-line arguments. With this update, a short usage help message is printed in the described scenario, describing the available options.

BZ#725497

When the tuned utility was running in a virtual guest, the disk-scheduler setting was not applied on virtual disks (vd*). Now, default configuration files that cover virtual disks have been updated and the disk-scheduler setting is now applied on the virtual disks in virtual guests.

Enhancements

BZ#740976

With this update, a new "virtual-host" profile has been added to the tuned package, providing fine-tuned profile for hypervisors managed by Red Hat Enterprise Virtualization Manager.

BZ#740977

With this update, a new "virtual-guest" profile for virtual systems has been added to the tuned package, providing fine-tuned profile for Red Hat Enterprise Linux 6 KVM virtual guests.

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

5.338. TZDATA

5.338.1. RHEA-2012:1488 – tzdata enhancement update

A new tzdata package that updates Daylight Saving Time observations for several countries is now available.

The tzdata packages contain data files with rules for various time zones around the world.

This updated package adds the following time-zone changes to the zone info database:

Bug Fix

BZ#871993, 871791, 871994, 871995

On October 24 2012, the Jordanian Cabinet rescinded a 2012-10-14 instruction to switch from daylight saving time (DST) to standard time on 2012-10-26. Instead, Jordan will remain on local DST (ITC +3) for the 2012-2013 Jordanian winter.

Cuba, which was scheduled to move back to standard time on 2012-11-12, switched to standard time on 2012-11-04.

BZ#871993, 871791, 871994, 871995

In Brazil, the North Region state, Tocantins, will observe DST in 2012-2013. This is the first time Tocantins has observed DST since 2003. By contrast, Bahia, a Northeast Region state, will not observe DST in 2012-2013. Like Tocantins, Bahia stopped observing DST in 2003. Bahia re-introduced DST on October 16 2011. On October 17 2012, however, Bahia Governor, Jaques Wagner, announced DST would not be observed in 2012, citing public surveys showing most Bahia residents were opposed to it.

BZ#871993, 871791, 871994, 871995

Israel has new DST rules as of 2013. DST now starts at 02:00 on the Friday before the last Sunday in March. DST now ends at 02:00 on the first Sunday after October 1, unless this day is also the second day of (Rosh Hashanah). In this case, DST ends a day later, at 02:00 on the first Monday after October 2.

The Palestinian territories, which were scheduled to move back to standard time on 2012-09-28, switched to standard time on 2012-09-21.

Although Western Samoa has observed DST for two consecutive seasons (2010-2011 and 2011-2012), there is no official indication of DST continuing according to a set pattern for the foreseeable future. On 2012-09-04, the Samoan Ministry of Commerce, Industry, and Labour announced Samoa would observe DST from Sunday, 2012-09-30 until Sunday 2012-04-07.

All users, especially those in the locale affected by these time changes, and users interacting with people or systems in the affected locale, are advised to upgrade to this updated package, which includes these updates.

5.338.2. RHEA-2012:1101 — tzdata enhancement update

Updated tzdata packages that add one enhancement are now available for Red Hat Enterprise Linux.

The tzdata packages contain data files with rules for various time zones around the world.

Enhancement

BZ#839271, BZ#839934, BZ#839937, BZ#839938

Daylight Saving Time will be interrupted during the holy month of Ramadan in Morocco (that is July 20 - August 19, 2012 in the Gregorian Calendar). This update incorporates the exception so that Daylight Saving Time is turned off and the time setting returned back to the standard time during Ramadan.

All users of tzdata are advised to upgrade to these updated packages, which add this enhancement.

5.338.3. RHEA-2013:0182 — tzdata enhancement update

New tzdata packages that add one enhancement are now available for Red Hat Enterprise Linux 3, 4, 5, and 6.

The tzdata packages contain data files with rules for time zones.

Enhancement

BZ#894030, BZ#894044, BZ#894045, BZ#894046

On Nov 10, 2012, Libya changed to the time zone UTC+1. Therefore, starting from the year 2013 Libya will be switching to daylight saving time on the last Friday of March and back to the standard time on the last Friday of October. The time zone setting and the daylight saving time settings for the Africa/Tripoli time zone have been updated accordingly.

All users of tzdata are advised to upgrade to these updated packages, which add this enhancement.

5.338.4. RHEA-2012:1338 — tzdata enhancement update

Updated tzdata packages that add two enhancements are now available for Red Hat Enterprise Linux.

The tzdata packages contain data files with rules for various time zones around the world.

Enhancements

BZ#857904, BZ#857905, BZ#857906, BZ#857907

Daylight saving time in Fiji will start at 2:00 a.m. on Sunday, 21st October 2012, and end at 3 am on Sunday, 20th January 2013.

BZ#857904, BZ#857905, BZ#857906, BZ#857907

Tokelau was listed in an incorrect time zone for as long as the Zoneinfo project was in existence. The actual zone was supposed to be GMT-11 hours before Tokelau was moved to the other side of the International Date Line at the end of year 2011. The local time in Tokelau is now GMT+13.

All users of tzdata are advised to upgrade to these updated packages, which add these enhancements.

5.338.5. RHEA-2013:0674 — tzdata enhancement update

Updated tzdata packages that add one enhancement are now available for Red Hat Enterprise Linux 3, 4, 5 and 6.

The tzdata packages contain data files with rules for various time zones.

Enhancement

BZ#921173, BZ#921174, BZ#919628, BZ#921176

Time zone rules of tzdata have been modified to reflect the following changes:

The period of Daylight Saving Time (DST) in Paraguay will end on March 24 instead of April 14.

Haiti will use US daylight-saving rules in the year 2013.

Morocco does not observe DST during Ramadan. Therefore, Morocco is expected to switch to Western European Time (WET) on July 9 and resume again to Western European Summer Time (WEST) on August 8.

Also, the tzdata packages now provide rules for several new time zones: Asia/Khandyga, Asia/Ust-Nera, and Europe/Busingen.

All users of tzdata are advised to upgrade to these updated packages, which add this enhancement.

5.338.6. RHEA-2013:1432 — tzdata enhancement update

Updated tzdata packages that add one enhancement are now available for Red Hat Enterprise Linux 3, 4, 5, and 6.

The tzdata packages contain data files with rules for various time zones.

Enhancement

BZ#1013527, BZ#1013875, BZ#1013876, BZ#1014720

Morocco extended DST by one month requiring an update to these packages. This update includes resynchronization with the latest upstream release in order to pick up the Moroccan DST change.

All users of tzdata are advised to upgrade to these updated packages, which add this enhancement.

5.338.7. RHEA-2013:0880 — tzdata enhancement update

Updated tzdata packages that add various enhancements are now available for Red Hat Enterprise Linux 3, 4, 5, and 6.

The tzdata packages contain data files with rules for various time zones.

Enhancement

BZ#928461, BZ#928462, BZ#928463, BZ#928464

The Gaza Strip and the West Bank entered Daylight Saving Time on March 28 at midnight local time.

All users of tzdata are advised to upgrade to these updated packages, which add these enhancements.

5.338.8. RHEA-2013:1025 – tzdata enhancement update

Updated tzdata packages that add one enhancement are now available for Red Hat Enterprise Linux 3, 4, 5, and 6.

The tzdata packages contain data files with rules for various time zones.

Enhancement

BZ#980805, BZ#980807, BZ#981019, BZ#981020

Morocco does not observe DST during Ramadan. Therefore, Morocco is expected to switch to Western European Time (WET) on July 7 and resume again to Western European Summer Time (WEST) on August 10. Also, the period of DST in Israel has been extended until the last Sunday in October from the year 2013 onwards.

All users of tzdata are advised to upgrade to these updated packages, which add this enhancement.

5.338.9. RHEA-2013:0615 – tzdata enhancement update

Updated tzdata packages that add one enhancement are now available for Red Hat Enterprise Linux 3, 4, 5 and 6.

The tzdata packages contain data files with rules for various time zones.

Enhancement

BZ#912521, BZ#916272, BZ#916273, BZ#916274

The Chilean Government is extending the period of Daylight Saving Time (DST) in the year 2013 until April the 27th. Then, Chile Standard Time (CLT) and Easter Island Standard Time (EAST) will be in effect until September the 7th when switching again to DST. With this update, the rules used for Chile time zones have been adjusted accordingly.

All users of tzdata are advised to upgrade to these updated packages, which add this enhancement.

5.339. UDEV

5.339.1. RHBA-2012:1007 – udev bug fix update

Updated udev packages that fix one bug 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.

Bug Fix

BZ#829703

Due to a bug in the binutils linker, the libudev library lost the ExecShield (GNU_RELRO) section, and was no longer protected by the Exec Shield security mechanism. This update provides a patch which ensures that the libudev library contains the ExecShield (GNU_RELRO) section again.

All users of udev are advised to upgrade to these updated packages, which fix this bug.

5.339.2. RHBA-2012:0906 – udev bug fix and enhancement update

Updated udev packages that fix multiple bugs 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 package replaces the devfs package and provides better hot plug functionality.

Bug Fixes

BZ#784648

Previously, libudev's function udev_device_get_devnode() returned NULL if the function was called before the udevd daemon processed the uevent for the device node. With this update, the udev_device_get_devnode() function now returns the devnode provided by the kernel.

BZ#735410

Previously, the udev-post non-service script located in the /etc/init.d/ directory always returned a failure code on a status request. This bug has been fixed in this update so that the udev-post service now returns information on whether it has already been run.

BZ#628762

Previously, the udev(7) man page did not mention the "nowatch" option. With this update, the "nowatch" option is now properly documented in the udev(7) man page.

All users of udev are advised to upgrade to these updated packages, which fix these bugs.

5.340. UNIXODBC

5.340.1. RHBA-2012:1509 – unixODBC bug fix update

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

The unixODBC packages contain a framework that supports accessing databases through the ODBC protocol.

Bug Fix

BZ#876488

When the isql utility was running in batch mode, which was activated by the "-b" command-line option, isql terminated unexpectedly with a segmentation fault upon receiving any SQL query. This bug has been fixed and isql no longer crashes in the described scenario.

All users of unixODBC are advised to upgrade to these updated packages, which fix this bug.

5.341. UPSTART

5.341.1. RHBA-2012:0863 – upstart bug fix and enhancement update

An updated upstart package that fixes two bug and adds two enhancements is now available for Red Hat Enterprise Linux 6.

The upstart package contains an event-based replacement for the /sbin/init daemon that starts tasks and services during boot, stops them during shut down, and supervises them while the system is running.

Bug Fixes

BZ#771736

Previously, the PACKAGE_BUGREPORT variable pointed to a Ubuntu mailing list. The mailing list was therefore presented in multiple manual pages, which was unwanted. With this update, the value of the PACKAGE_BUGREPORT variable has been modified to "https://launchpad.net/upstart/+bugs", and users are now directed to that website rather than to the Ubuntu mailing list.

BZ#798551

Previous versions of upstart did not mount the proc and sys file systems. This was ensured by initscripts, which could, under certain circumstances, lead to race condition problems. With this update, upstart is used to mount the proc and sys file systems before launching anything else.

Enhancements

BZ#663594

Files with the ".conf" suffix located in the /etc/init/ directory are not considered as configuration files. As a consequence, such files are not protected during a package update and can be overwritten by new files. This update adds support for "override" files that contain user-specified settings. Now, it is possible to alter parameters provided by the aforementioned ".conf" files by creating a corresponding file with the ".override" suffix.

BZ#735427

Previously, the initctl scripts returned error messages that did not tell users how to run the particular command correctly to get the required output. This update adds a new stanza, "usage", that can be used to provide users with detailed information on how to run the particular command correctly if the input has been incorrect.

All users of upstart are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.342. USBREDIR

5.342.1. RHBA-2012:1435 – usbredir bug fix update

Updated usbredir packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The usbredir packages provide a network protocol for sending USB device traffic over a network connection and a number of libraries to help implement support for this protocol.

Bug Fix

BZ#858776

Due to a bug in the libusbredirhost library, handling of timeouts for bulk transfers did not work correctly. Consequently, traffic of USB ACM serial port devices, such as PSTN modems and SmartCard readers, could not be properly redirected. With this update, no timeout is set on the usb-host side for these devices and the traffic redirection works as expected.

Users who use USB redirection for Spice are advised to upgrade to these updated packages, which fix this bug.

5.343. UTIL-LINUX-NG

5.343.1. RHBA-2012:1427 – util-linux-ng bug fix update

Updated util-linux-ng packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The util-linux-ng packages contain a set of low-level system utilities that are necessary for a Linux operating system to function.

Bug Fix

BZ#864367

When the telnetd daemon was used to log in to a server, the login utility failed to update the /var/run/utmp file properly. Consequently, the line used for a previous session in /var/run/utmp was not reused, thus growing the file unnecessarily. A patch has been provided to address this issue and the login utility now always updates /var/run/utmp as expected.

Users of util-linux-ng are advised to upgrade to these updated packages, which fix this bug.

5.343.2. RHBA-2012:0925 – util-linux-ng bug fix update

Updated util-linux-ng packages that fix several bugs are now available for Red Hat Enterprise Linux 6.

The util-linux-ng packages contain a set of low-level system utilities that are necessary for a Linux operating system to function.

Bug Fixes

BZ#588419

The console login time-out value was set to 60 seconds. This could cause the login to time out during the name lookup process on systems with broken DNS (Domain Name Service). With this update, the timeout value has been prolonged to 180 seconds to allow the login process to complete name lookups under these circumstances.

BZ#740163

The "fdisk -l" and "sfdisk -l" commands returned confusing warnings for unpartitioned devices similar to the following:

> Disk /dev/mapper/[volume name] doesn't contain a valid partition table

With this update, the commands ignore unpartitioned devices and the problem no longer occurs.

BZ#785142

Previously, after the installation of the uuidd package, the uuidd daemon was not enabled by default. With this update, the underlying code has been modified and the uuidd daemon is enabled after installation as expected and can be started by the init script after reboot.

BZ#797888

Previously, the script command did not work correctly if called from the csh shell in the /etc/csh.login file. The child processes created by the script inherited the SIGTERM ignore property from csh and could not be terminated with the signal. With this update, the script resets the SIGTERM setting so that the shell is started with the default SIGTERM behavior and its children accept signals as expected.

All users of util-linux-ng are advised to upgrade to these updated packages, which fix these bugs.

5.344. VALGRIND

5.344.1. RHBA-2012:0936 – valgrind bug fix and enhancement update

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

The valgrind packages provide a programming utility for debugging memory, detecting memory leaks, and profiling.

Bug Fix

BZ#757728

Prior to this update, the "memalign" and "posix_memalign" replacements could only handle alignments of 1 MB maximum. As a consequence, running qemu-kvm in valgrind could cause alignment errors. This update modifies the underlying code so that memalign and posix_memalign replacement can now handle alignments up to 4 MB.

Enhancement

BZ#739143

With this update, valgrind has been updated to provide complete support for IBM POWER7 Series and VPN-1 Power VSX hardware as well as support for Decimal Floating Point (DFP).

All users of valgrind are advised to upgrade to these updated packages, which fix this bug add this enhancement.

5.345. VIM

5.345.1. RHBA-2012:0454 – vim bug fix update

Updated vim packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

Vim (Vi IMproved) is an updated and improved version of the vi editor.

Bug Fixes

BZ#594997

Previously, when using the VimExplorer file manager with the locale set to Simplified Chinese (zh_CN), the netrw.vim script inserted an unwanted "e" character in front of file names. The underlying code has been modified so that file names are now displayed correctly, without unwanted characters.

BZ#634902

The spec file template that was used when new spec files were edited contained outdated information. With this update, the spec file template is updated to adhere to the latest spec file guidelines.

BZ#652610

When using the file explorer in a subdirectory of the root directory, the "vim .." command displayed only part of the root directory's content. A patch has been applied to address this issue, and the "vim .." command now lists the content of the root directory properly in the described scenario.

BZ#663753

Due to a typographic error in the filetype plug-in, the vim utility could display the httpd configuration files with incorrect syntax highlighting. This update corrects the errors in the filetype plug-in, and the httpd configuration files are now displayed with the correct syntax highlighting.

All users of vim are advised to upgrade to these updated packages, which fix these bugs.

5.346. VINO

5.346.1. RHSA-2013:0169 – Moderate: vino security update

An updated vino package that fixes several security issues is 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 score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) associated with each description below.

Vino is a Virtual Network Computing (VNC) server for GNOME. It allows remote users to connect to a running GNOME session using VNC.

Security Fixes

CVE-2012-4429

It was found that Vino transmitted all clipboard activity on the system running Vino to all clients connected to port 5900, even those who had not authenticated. A remote attacker who is able to access port 5900 on a system running Vino could use this flaw to read clipboard data without authenticating.

CVE-2011-0904, CVE-2011-0905

Two out-of-bounds memory read flaws were found in the way Vino processed client framebuffer requests in certain encodings. An authenticated client could use these flaws to send a specially-crafted request to Vino, causing it to crash.

CVE-2011-1164

In certain circumstances, the vino-preferences dialog box incorrectly indicated that Vino was only accessible from the local network. This could confuse a user into believing connections from external networks are not allowed (even when they are allowed). With this update, vino-preferences no longer displays connectivity and reachable information.

CVE-2011-1165

There was no warning that Universal Plug and Play (UPnP) was used to open ports on a user's network router when the "Configure network automatically to accept connections" option was enabled (it is disabled by default) in the Vino preferences. This update changes the option's description to avoid the risk of a UPnP router configuration change without the user's consent.

All Vino users should upgrade to this updated package, which contains backported patches to resolve these issues. The GNOME session must be restarted (log out, then log back in) for this update to take effect.

5.347. VIOS-PROXY

5.347.1. RHBA-2012:0755 – vios-proxy bug fix update

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

The vios-proxy program suite creates a network tunnel between a server in the QEMU host and a client in a QEMU guest. The proxied server and client programs open normal TCP network ports on localhost and the vios-proxy tunnel connects them using QEMU virtioserial channels.

Bug Fix

BZ#743723

Previously, the packages did not contain manual pages for the vios-proxy-host and vios-proxy-guest daemons. With this update, these manual pages are now available.

All users of vios-proxy are advised to upgrade to these updated packages, which fix this bug.

5.348. VIRTIO-WIN

5.348.1. RHBA-2012:1083 – virtio-win bug fix update

An updated virtio-win package that fixes two bugs is 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#838523

A bug in the virtio serial driver could cause a Stop Error (also known as Blue Screen of Death, or BSoD) which occurred on a guest machine when transferring data from the host. This update fixes the bug in the driver so that the guest machine no longer crashes with Blue Screen of Death in this scenario.

BZ#838655

The QXL driver included in the previous version of the virtio-win package was not digitally signed. The QXL driver provided in this update in digitally signed.

All users of virtio-win are advised to upgrade to this updated package, which fixes these bugs.

5.348.2. RHBA-2012:0751 – virtio-win bug fix and enhancement update

An updated virtio-win package that fixes multiple bugs and adds various enhancements is 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#492777

Previously, if a Microsoft Windows guest machine sent more tx fragments than 256 (the ring size), the NetKVM driver dropped packets. To prevent this problem, indirect ring support has been implemented in the NetKVM driver.

BZ#759361

Previously, user were not able to update the rx and tx parameters in Windows Registry by using the NetKVMConfig utility. Although the utility reported that the parameters had been changed, the change was not displayed in the Windows Device Manager. This was due to incorrect NetKVMConfig parameters changing handler, which has been fixed, so that NetKVMConfig now works as expected and users can update the rx and tx parameters.

BZ#753723

Previously, the block driver (viostor) did not provide support for obtaining serial numbers of virtio block devices from QEMU. The serial numbers were therefore not available on Windows guest machines. With this update, a serial number of a virtio block device is now retrieved from miniport during the find adapter phase.

BZ#752743

Prior to this update, the block driver (viostor) driver did not reject write requests to read-only volumes. Attempting to format a read-only volume caused the guest to stop with an EIO error. With this update, if the target volume has the read-only flag, the guest does not stop, and write requests are completed with an error. Attempting to format or write to a read-only volume are now rejected by the viostor driver.

BZ#751952

Previously, if the "Fix IP checksum on LSO" option in Microsoft Windows Device Manager was disabled, users were not able to transfer data from a guest machine to the host machine using the winscp utility. To prevent this problem, it is no longer possible to disable the "Fix IP checksum on LSO" option.

BZ#803950

A bug in the balloon driver could cause a stop error (also known as Blue Screen of Death, or BSoD) if a guest machine entered the S3 (suspend to RAM) or S4 (suspend to disk) state while performing memory balooning on it. The bug in the balloon driver has been fixed, and the stop error no longer occurs under these circumstances.

BZ#810694

Previously, incorrect flush requests handling could lead to a race condition in the block driver (viostor). Under heavy load, usually when using the "cache=writeback" option, the flush handler was executed asynchronously without proper synchronization with the rest of request processing logic. With this update, execution of the flush request is synchronized with the virtio Interrupt Service Routine (ISR), and the race condition no longer occurs in this scenario.

BZ#771390

The viostor utility did not check the size of an incoming buffer. Applications could send buffers larger than the maximum transfer size to the viostor driver directly by bypassing the file system stack. The buffer size is now reduced if it is bigger that the maximum transfer size. The viostor driver can now properly handle requests with buffers of any size.

Enhancements

BZ#677219

Previously, it was not possible to resize non-system disks online, without reboot. This update adds support for online resizing of VirtIO non-system disks.

BZ#713643

This update provides optimized offload RX IP checksum for the virtio_net driver.

BZ#808322

Offload parameters for the virtio-win network driver have been updated. Multiple parameters are now set to "enabled" by default. To edit parameters of an installed driver, open Microsoft Windows Device Manager, choose "Red Hat VirtIO Ethernet Adapter" from the "Network Adapters" list and click on the "Advanced" tab.

All users of virtio-win are advised to upgrade to this updated package, which fixes these bugs and adds these enhancements.

5.349. VIRT-MANAGER

5.349.1. RHBA-2012:0785 – virt-manager bug fix and enhancement update

Updated virt-manager packages that fix multiple bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The Virtual Machine Manager (virt-manager) is a graphical tool for administering Kernel-based Virtual Machines (KVM) using libvirt.

Bug Fixes

BZ#741937, BZ#769192

Prior to this update, the virt-manager tool did not change all required XML elements when changing the disk bus of a guest from IDE to virtio, or vice versa. As a consequence, virt-manager could create invalid XML elements that libvirt accepted but then the guest failed to boot. This update modifies the underlying code to change all required XML elements and guests start now as expected.

BZ#742055

Prior to this update, the virt-manager tool could, under certain circumstances, receive an error from libvirt if virt-manager tried to read a domain's information while that domain was shutting down. As a consequence, the libvirt connection in the user interface (UI) was incorrectly closed. This update modifies the underlying code to expect errors in this case and not close the libvirt connection.

BZ#747490

Prior to this update, the glib integration did not work correctly. As a consequence, serial consoles of virtual machines could stall when transferring large amounts of data. This update modifies the underlying code to allow the transfer of larger amounts of data without stalling.

BZ#750225

Prior to this update, graphical scaling did not work for SPICE graphics in the virt-manager tool. This update modifies the underlying code to connect the UI scaling selection element to the spice back end.

BZ#803600

Prior to this update, the virt-manager tool could close with a segmentation fault if the user deleted several storage volumes in quick succession due to data locking when deleting the storage volume. This update modifies the virt-manager threads to allow deleting of storage volumes in quick succession.

BZ#811316

Prior to this update, the virt-manager tool could, did not correctly clean up certain internal state. As a consequence, closing and reopening a graphical window of a guest did not reopen the graphical console. This update modifies the underlying code to ensure that the graphical console connection is correctly reopened.

BZ#816279

Prior to this update, the virt-manager tool did not correctly using the graphics listen address attribute in the configuration of the virtual machine, and would always try to connect to the guest with an SSH tunnel. This could break graphical console connections if "listen=" was set to an explicit interface address. This update modifies the underlying code to ensure that virt-manager now correctly connects to these addresses.

Enhancement

BZ#716673

Prior to this update, the default disk image format for disk images used for newly created virtual machines was "raw" and could not be changed. This update adds the new "qcow2" option.

All users of virt-manager are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.350. VIRT-TOP AND OCAML-LIBVIRT

5.350.1. RHBA-2012:0959 – virt-top and ocaml-libvirt bug fix update

Updated virt-top and ocaml-libvirt packages that fix two bugs are now available.

The virt-top utility displays statistics of virtualized domains and uses many of the same keys and command line options as the top utility.

The ocaml-libvirt package provides OCaml (Objective CAML) bindings for libvirt, allowing users to write OCaml programs and scripts which control virtualization features.

Bug Fix

BZ#737728

Output of the "virt-top -1" command (which displays physical CPU usage) did not contain all the needed information from libvirt in order to provide accurate accounting of physical CPU usage. With this update, the underlying source code has been modified to address this issue, and the "virt-top -1" output now displays accurate statistics.

All users of virt-top and ocaml-libvirt are advised to upgrade to these updated packages, which resolve these issues.

5.351. VIRT-V2V

5.351.1. RHBA-2012:1468 – virt-v2v bug fix update

Updated virt-v2v packages that fix a bug are now available for Red Hat Enterprise Linux 6.

The virt-v2v packages provide 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#872496

Previously, when the virt-v2v utility was used to convert a virtual machine (VM) from a foreign hypervisor (such as Xen or VMware) to Red Hat Enterprise Virtualization, it set the vm_snapshot_id identifier of all disks of that VM incorrectly. Consequently, various problems occurred while doing a large number of tasks on this VM from the side of Red Hat Enterprise Virtualization. With this update, a unique identifier is generated for each disk in the described scenario, thus preventing this bug.

Users of virt-v2v are advised to upgrade to these updated packages, which fix this bug.

5.351.2. RHBA-2012:0788 – virt-v2v bug fix and enhancement update

Updated virt-v2v packages that fix multiple bugs and add various enhancements are now available for Red Hat Enterprise Linux 6.

The virt-v2v packages provide 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 packages have been upgraded to upstream version 0.8.7, which provides a number of bug fixes over the previous version. (BZ#753732)

Bug Fixes

BZ#737600

Previously, a converted Microsoft Windows XP guest could terminate unexpectedly on boot with a STOP error (also known as Blue Screen of Death, or BSoD). The error could be triggered if the guest was configured with a CPU or chipset driver that malfunctioned when the CPU or chipset was not present; this only occurred when converting Physical-to-Virtual, or P2V, machines. With this update, the registry keys related to certain services known to cause problems are deleted during the conversion process. The converted guest now boots as expected after the conversion process.

BZ#767262

When converting a Physical-to-Virtual (P2V) or a Virtual-to-Virtual (V2V) machine to run on Red Hat Enterprise Virtualization systems, **virt-v2v** failed with a write error when attempting to write to the target export storage domain. This happened if the target system did not use the standard UID and GID for ownership of the target Red Hat Enterprise Virtualization export storage domain. With this update, **virt-v2v** checks the local system for the UID of "vdsm" and the GID of "kvm". If present, the values are treated as the values required to write to the Red Hat Enterprise Virtualization export storage domain, and the conversion succeeds in the described scenario.

BZ#751293

When running the **virt-v2v** utility and the **/var/lib/virt-v2v/** directory did not contain any files other than **virt-v2v.db**, the conversion failed with an error message similar to:

```
/transfer0w34SV: umount: /sysroot/transfer0w34SV: not mounted at
/usr/share/perl5/vendor_perl/Sys/VirtConvert/GuestfsHandle.pm line 193.
at /usr/share/perl5/vendor_perl/Sys/VirtConvert/Config.pm line 262
```

The underlying source code has been modified to correctly handle situations when there is no software available locally for installation into a guest during conversion, and so ensures that the conversion succeeds.

BZ#737855

When converting a Xen HVM guest with references to **/dev/xvdX** devices in the fstab or GRUB device map file, the **/dev/xvdX** devices in these files were not updated. With this update, the virt-v2v and virt-p2v utilities now look in the Xen HVM guest configuration files during the conversion process for devices named **/dev/xvdX** as well as **/dev/hdX**. Both are treated as identical and either is converted to **/dev/vdX**. References to Xen paravirtualized block devices in fstab and device map of Xen HVM guest are now correctly updated during the conversion.

BZ#696779

Previously, the **virt-v2v** utility unconditionally marked all converted guests as a Server-type workload. This caused Desktop-type workload guests to be displayed incorrectly in **Red Hat Enterprise Virtualization Manager**. This update adds a new command-line option, **--vmtype**, which forces the conversion process to mark the newly created Red Hat Enterprise Virtualization virtual machine as either **Desktop** or **Server**. If **--vmtype** is omitted, virt-v2v attempts to determine the correct type.

BZ#787734

The new VMware Tools are split into multiple packages. Previously, when converting such a guest, the VMware Tools packages were not removed during the conversion process. This could cause warnings to be displayed in the converted guest or cause the guest to function incorrectly. The conversion process has been updated to recognize the new VMware Tools packages and remove them. The VMware Tools packages are now correctly removed during the conversion.

BZ#786115

When attempting to convert a guest which was accessed over an SSH connection and the target host had an SSH login banner configured, the conversion process could become unresponsive. With this update, SSH login banners are ignored and the conversion process completes as expected.

Enhancements

BZ#695406

Previously, the virt-v2v utility could be used to move a virtual machine from one environment to another, but not to move a workload from a physical server. With this update, users can move the server data over the network to a virtual environment by using the new **virt-p2v** tool.

BZ#768172

If the user used third-party kernel modules in a guest machine, and updated the kernel, the conversion could fail or the converted guest could fail to operate correctly. This was because the conversion process did not recognize third-party kernel modules. Users can now specify a "user-custom" capability for the guest operating system in the **virt-v2v.conf** file. All the dependencies of "user-custom" are installed during the conversion process.

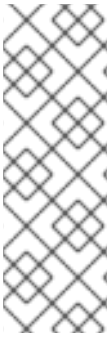
All users of virt-v2v are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.352. VIRT-VIEWER

5.352.1. RHBA-2012:0772 – virt-viewer bug fix and enhancement update

Updated virt-viewer packages that fix multiple 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 and SPICE clients.



NOTE

The virt-viewer packages have been upgraded to upstream version 0.5.2, which provides a number of enhancements over the previous version. Among these, this update adds support for a new client, known as **remote-viewer**, which obsoletes the need for the separately maintained **spicec** application. The **remote-viewer** utility provides a way for the Simple Protocol for Independent Computing Environments (SPICE) XPI plugin to connect to remote SPICE servers, which ensures a consistent user experience of the virt-viewer utility. (BZ#784920)

Bug Fixes

BZ#749759

The SPICE client cannot determine in advance whether a SPICE server requires a password for authentication. Thus, the SPICE client attempts to establish the connection and if it receives an authentication error, the client closes the connection, prompts the user for a password, and then reconnects to the SPICE server. Previously, the SPICE client was unable to perform the reconnection step when the connection to the server was using SSH tunneling. With this update, the SPICE client is allowed to request a new SSH tunnel connection to the SPICE server after obtaining the password from the user. Now, users are able to connect to password-protected SPICE servers when using an SSH tunnel.

BZ#784922

The **virt-viewer** utility has been modified to enable support for USB redirection introduced in the latest version of the spice-gtk packages. Users are now able to attach local USB devices to remote virtual machines using the **SPICE** protocol.

BZ#811191

Previously, the virt-viewer manual page did not describe the **--attach**, or **-a**, option. With this update, the virt-viewer manual page explains that **libvirt** can be used to directly attach **virt-viewer** to a local display instead of making a TCP/UNIX socket connection when using one of the aforementioned options.

BZ#749723

When running a guest while the user password is set, the **virt-viewer** application asks for authentication. However, due to incorrect signal handling, if the user canceled the dialog box, the following error message was returned:

```
Unable to authenticate with remote desktop server at localhost:5900:
Unable to collect credentials.Retry connection again?
```

The underlying source code has been modified to ensure correct signal handling. Now, if **virt-viewer** receives a signal about a session being canceled, **virt-viewer** is disconnected and exits without error messages, as expected.

BZ#813375

Previously, the URI parsing code did not expect URIs containing square brackets, **[** and **]**, around the host component. It was thus not possible to connect to a remote libvirt server whose URI address contained raw IPv6 addresses (for example **qemu+ssh://root@[2001::xxxx:1]/system**). With this update, the URI parsing has been fixed to take account of the IPv6 address syntax, so it is now possible to connect to remote libvirt servers using raw IPv6 addresses.

BZ#810544

On 32-bit Intel architectures, an arithmetic error caused inaccurate calculation of the desired window size. The error manifested itself as a one-pixel black bar appended to the bottom of the window in full-screen mode, thus causing the guest display to be unnecessarily scaled. The scaling code has been changed to round to the nearest integer instead of truncating, which avoids a reliance on precision of floating point calculations. On 32-bit Intel architectures, windows are resized such that scaling is not required if the guest display is small enough to fit on the host desktop.

BZ#819436

Due to a race condition, the following message could be displayed at the command line when closing the **virt-viewer** application:

```
Segmentation fault (core dumped)
```

The underlying source code has been modified to prevent the race condition from occurring, and virt-viewer now exits gracefully, without error messages.

BZ#816550

When reconnecting a guest with multiple monitors (for example after a restart), **virt-viewer** created new windows for the additional monitors, while the old windows still existed. This was because the **GtkWindow** object was not freed. This update modifies **virt-viewer** ensure that windows are closed when a display closes.

BZ#816280

With this update, the **OK** button label of the **USB device selection** dialog box has been changed to **Close**.

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

5.353. VIRT-WHO

5.353.1. RHBA-2012:0900 – virt-who bug fix and enhancement update

Updated virt-who packages that fix several bugs and add various enhancement are now available.

The virt-who package provides an agent that collects information about virtual guests present in the system and reports them to the subscription manager.

The virt-who package has been upgraded to upstream version 0.6, which provides a number of bug fixes and enhancements over the previous version. (BZ#790000)

Bug Fixes

BZ#746163

When the libvirtd daemon was stopped, the virt-who utility no longer received information about the state of the guest, and showed an inaccurate list of guest UUIDs. With this update, polling is used to check the connection to the libvirtd daemon, and the time for which the list of UUIDs is inaccurate was minimized.

BZ#813299

Prior to this update, the virt-who utility could not connect to the libvirt daemon due to a regression in the code that handles forking of the virt-who daemon. With this update, a connection is open to the libvirtd daemon after the fork of the virt-who daemon; thus, fixing this issue.

BZ#801657

This update includes a missing python-suds dependency into the virt-who specfile. The missing dependency was causing the virt-who daemon to fail to start.

BZ#806225

The virt-who daemon did not use double forking when it was started. Consequently, the daemon did not detach from the terminal correctly. With this update, virt-who uses double forking, and is able to correctly detach itself from the terminal.

BZ#815279

Previously, the virt-who utility could not handle all events that were being sent by the libvirtd daemon. If an unrecognized event was received, virt-who logged an IndexError in the logs, and returned a traceback error. With this update, virt-who handles all error events (even unknown), and no longer returns a traceback error.

Users are advised to upgrade to these updated virt-who packages, which resolve these issues and add these enhancement.

5.354. VSFTPD

5.354.1. RHBA-2013:0263 – vsftpd bug fix update

An updated vsftpd package that fixes one bug is now available for Red Hat Enterprise Linux 6.

The vsftpd package includes a Very Secure File Transfer Protocol (FTP) daemon, which is used to serve files over a network.

Bug Fix

BZ#910371

The vsftpd daemon supports FTP clients that provide the set of commands "proxy ftp-command". These commands provide the ability to transfer data from one server to another through FTP client. Previously, the vsftpd version failed to establish data connections to another server opened with the "proxy get [file]" command and sent the data connection request to the client instead. With this update, the vsftpd version is able to establish data connections to another FTP server using the "proxy get [file]" command.

Users of vsftpd are advised to upgrade to this updated package, which fixes this bug. The vsftpd daemon must be restarted for this update to take effect.

5.354.2. RHBA-2012:0790 – vsftpd bug fix update

Updated vsftpd packages that fix multiple bugs are now available for Red Hat Enterprise Linux 6.

The vsftpd package provides the VSFTP (Very Secure File Transfer Protocol) daemon.

Bug Fixes

BZ#701300

Prior to this update, the configuration file specified the wrong default log file. As a consequence, the logrotate script could not find and consequently rotate the vsftpd log file which resulted in an unnecessarily large vsftpd log. This update specifies /var/log/xferlog as its default log file in /etc/vsftpd/vsftpd.conf, which enables log rotation on vsftpd log files.

BZ#708657

Prior to this update, the RLIMIT_AS value (100 MB) was insufficient. As a consequence, LDAP could not use vsftpd for authentication to the system. This update increases the initial RLIMIT_AS value to 200 MB, and vsftpd now can be used for LDAP authentication as expected.

BZ#717411

Prior to this update, vsftpd did not handle file transfer failures correctly if the ftp-data port was blocked on the File Transfer Protocol (FTP) client. As a consequence, vsftpd could become unresponsive. This update modifies the underlying code so that the vsftp daemon reports such failures to the FTP client and the data transfer is now terminated as expected.

BZ#745133

Prior to this update, the man page of the vsftpd.conf file contained incorrect default values for "max_per_ip" and "max_clients" options. This update introduces the correct default values for these two options.

BZ#752954

Prior to this update, the DNS reverse lookup feature could not be disabled. This update adds the "reverse_lookup_enable" parameter, which allows to enable or disable the DNS reverse lookup functionality.

BZ#765757

Prior to this update, vsftpd also listed the CHMOD command when the "chmod_enable" option was disabled. This update modifies the help file so that vsftpd no longer lists the CHMOD command when the command is disabled.

BZ#785061

Prior to this update, listing files could cause an overflow error if a directory contained files with a User or Group ID that was higher then the maximum value 2147483647 of the "signed int" data type. As a consequence, the FTP connection was terminated. This update modifies vsftpd to support UIDs and GIDs above the maximum value of the "unsigned int" data type. Directory content is now listed as expected in the scenario described.

BZ#785084

Prior to this update, the ls command did not support square brackets as wildcard characters in FTP connections. This update improves wildcard characters support in vsftpd and square brackets can now be used in regular expressions with the ls command.

BZ#785642

Prior to this update, the "listen()" function in vsftpd could, under certain circumstances, fail under heavy load. As a consequence, the socket became blocked. This update closes failed sockets and creates new a socket to cointinue listening.

All users of vsftpd are advised to upgrade to these updated packages, which fix these bugs.

5.355. WGET

5.355.1. RHBA-2012:1353 – wget bug fix update

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

The wget packages provide the GNU Wget file retrieval utility for HTTP, HTTPS, and FTP protocols. Wget provides various useful features, such as the ability to work in the background while the user is logged out, recursive retrieval of directories, file name wildcard matching or updating files in dependency on file timestamp comparison.

Bug Fixes

BZ#754168

Prior to this update, the wget package contained a redundant URL to the wget upstream project. This update modifies the specification file to list the correct http://www.gnu.org/software/wget/.

BZ#814208

Prior to this update, the wget utility did not previously work as intended with the "-T, --timeout" option set when http server did not answer the SSL handshake. Wget source code has been patched, to ensure that wget aborts the connection when using --timeout option correctly.

BZ#714893

Prior to this update, the wget utility source code was lacking check of the HTTP response parsing function return value. In some cases, when HTTP response header was malformed (fuzzed), the parsing function returned error. Because the returned value was not checked, it then resulted in Segmentation Fault. This update adds check of the HTTP response parsing function return value in the wget source code. Now when HTTP response header is malformed (fuzzed) and the parsing function returns error, the following error message is thrown and wget retries the request.

```
2012-10-01 10:13:44 ERROR -1: Malformed status line.
```

All users of wget are advised to upgrade to these updated packages, which fix this bug.

5.356. WORDNET

5.356.1. RHBA-2012:0932 – wordnet bug fix update

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

WordNet provides a set of utilities and a lexical database to manage English words in sets of synonyms (synsets). Wordnet uses these synsets to generate a combination of dictionary and thesaurus and to support automatic text analysis.

Bug Fix

BZ#658043

Prior to this update, WordNet encountered file conflicts when trying to install wordnet packages for 32-bit and 64-bit architectures on the same host. As a consequence, installing the second package failed. This update removes the conflicting auxiliary files from the packages. Now, WordNet is multi-architecture safe.

All users of wordnet are advised to upgrade to these updated packages, which fix this bug.

5.357. WPA_SUPPLICANT

5.357.1. RHBA-2013:0225 – wpa_supplicant bug fix and enhancement update

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

The wpa_supplicant packages contain an 802.1X Supplicant with support for WEP, WPA, WPA2 (IEEE 802.11i / RSN), and various EAP authentication methods. It implements key negotiation with a WPA Authenticator for client stations and controls the roaming and IEEE 802.11 authentication/association of the WLAN driver.

Bug Fix

BZ#855255

Previously, the supplicant would attempt to roam to slightly stronger access points, increasing the chance of a disconnection. This bug has been fixed and the supplicant now only attempts to roam to a stronger access point when the current signal is significantly degraded.

Enhancement

BZ#855273

Support for Opportunistic Key Caching (OKC), also known as Proactive Key Caching (PKC), has been added to WPA Supplicant to facilitate faster and less error-prone roaming between access points in the same network.

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

5.357.2. RHBA-2012:0978 – wpa_supplicant bug fix update

An updated wpa_supplicant package that fixes a bug is now available for Red Hat Enterprise Linux 6.

The wpa_supplicant package contains a WPA (Wi-Fi Protected Access) Supplicant utility for Linux, BSD, and Windows with support for WPA and WPA2 (IEEE 802.11i/RSN). The supplicant is an IEEE 802.1X/WPA component that is used in client workstations. It implements key negotiation with a WPA Authenticator and it controls the roaming and IEEE 802.11 authentication and association of the WLAN driver.

Bug Fix

BZ#752032

Due to an error in the wpa_supplicant code, Wi-Fi signal levels for some wireless devices reported by the Linux kernel's nl80211 API were not handled correctly. Consequently, the NetworkManager applet did not indicate the signal strength for unconnected networks. With this update, the code has been corrected and the NetworkManager applet now indicates signal strength for drivers using the nl80211 API as expected.

All users of wpa_supplicant are advised to upgrade to this updated package, which fixes this bug.

5.358. XFIG

5.358.1. RHBA-2012:0985 – xfig bug fix update

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

The xfig packages contain the Xfig editor. Xfig is an open-source vector graphics editor, which allows you to create simple diagrams and figures.

Bug Fix

BZ#806689

Security errata RHSA-2012:0095 changed the way ghostscript handles relative paths. Consequently, as Xfig relied on the original ghostscript behavior, it failed to open encapsulated postscript files and returned the execution stack ending with the following error:

```
Current allocation mode is local
Last OS error: 2
GPL Ghostscript 8.70: Unrecoverable error, exit code 1
EPS object read OK, but no preview bitmap found/generated
```

Xfig was changed to use absolute paths when executing the ghostscript binary. With this update, Xfig opens encapsulated postscript files and includes them in other figures as expected.

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

5.359. XFSPROGS

5.359.1. RHBA-2012:0883 – xfsprogs bug fix update

Updated xfsprogs packages that fix four 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 mkfs.xfs.

Bug Fixes

BZ#730886

Prior to this update, certain file names could cause the xfs_metadump utility to become suspended when generating obfuscated names. This update modifies the underlying code so that xfs_metadump now works as expected.

BZ#738279

Prior to this update, the allocation group size (agsize) was computed incorrectly during mkfs for some filesystem sizes. As a consequence, creating file systems could fail if file system blocks within an allocation group (agblocks) were increased past the maximum. This update modifies the computing method so that agblocks are no longer increased past the maximum.

BZ#749434

Prior to this update, the xfs_quota utility failed with the error message "xfs_quota: cannot initialise path table: No such file or directory" if an invalid xfs entry was encountered in the mtab. This update modifies the xfs_quota utility so that the xfs_quota utility now runs as expected.

BZ#749435

Prior to this update, the xfs_quota utility reported that the project quota values were twice as high as expected. This update modifies the xfs_quota utility so that it now reports the correct values.

All users who use the XFS file system are advised to upgrade to these updated packages, which fix these bugs.

5.360. XINETD

5.360.1. RHBA-2012:1162 – xinetd bug fix update

An updated xinetd package that fixes one bug is now available for Red Hat Enterprise Linux 6.

Xinetd is a secure replacement for inetd, the Internet services daemon. Xinetd provides access control for all services based on the address of the remote host and/or on time of access, and can prevent denial-of-access attacks. Xinetd provides extensive logging, has no limit on the number of server arguments, and allows users to bind specific services to specific IP addresses on a host machine. Each service has its own specific configuration file for Xinetd; the files are located in the /etc/xinetd.d directory.

Bug Fix

BZ#841916

Due to incorrect handling of a file descriptor array in the service.c source file, some of the descriptors remained open when xinetd was under heavy load. Additionally, the system log was filled with a large number of messages that took up a lot of disk space over time. This bug has been fixed in the code, xinetd now handles the file descriptors correctly and no longer fills the system log.

All users of xinetd are advised to upgrade to this updated package, which fixes this bug.

5.360.2. RHBA-2012:0409 – xinetd bug fix update

An updated xinetd package that fixes multiple bugs is now available for Red Hat Enterprise Linux 6.

The xinetd daemon is a secure replacement for xinetd, the Internet services daemon. The xinetd daemon provides access control for all services based on the address of the remote host, on time of access, or both, and can prevent denial of service (DoS) attacks.

Bug Fixes

BZ#694820

Under certain circumstances, the xinetd daemon could become unresponsive (for example, when trying to acquire an already acquired lock for writing to its log file) when an unexpected signal arrived. With this update, the daemon handles unexpected signals correctly and no longer hangs under these circumstances.

BZ#697783

Previously, a bug in the xinetd code could cause corruption of the time_t variable resulting in the following compiler warning:

```
warning: dereferencing type-punned pointer will break strict-aliasing rules
```

A patch has been applied to address this issue, so that the warning no longer occurs.

BZ#697788

Previously, the xinetd daemon ignored the "port" line of the service configuration file, and it was therefore impossible to bind certain RPC services to a specific port. The underlying source code has been modified to ensure that xinetd honors the "port" line, so that the port numbers are now handled appropriately.

BZ#711787

Incorrect use of the realloc() function could cause memory corruption. This resulted in the xinetd daemon terminating unexpectedly right after the start when a large number of services had been configured. The realloc() function has been removed, which ensures that memory corruption no longer occurs in this scenario, and the xinetd daemon starts successfully even when configuring a large number of services.

All users of xinetd are advised to upgrade to this updated package, which fixes these bugs.

5.361. XMLRPC-C

5.361.1. RHBA-2012:0954 – xmlrpc-c bug fix update

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

The xmlrpc-c packages provide a network protocol to allow a client program to make a simple RPC (remote procedure call) over the Internet. It converts an RPC into an XML document, sends it to a remote server using HTTP, and gets back the response in XML.

Bug Fixes

BZ#653702

Prior to this update, the "xmlrpc-c-config client --libs" command returned unprocessed output, making it difficult to discern important information from it. This bug has been fixed and the output of the command is now properly pre-processed by the autoconf utility.

BZ#741641

A memory leak was discovered in the xmlrpc-c library by the valgrind utility. A patch has been provided to address this bug and the memory leak no longer occurs.

Users of xmlrpc-c are advised to upgrade to these updated packages, which fix these bugs.

5.362. XORG-X11-DRV-ATI AND MESA

5.362.1. RHEA-2012:0903 – xorg-x11-drv-ati and mesa bug fix and enhancement update

Updated xorg-x11-drv-ati and mesa packages that fix a bug and add an enhancement 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. The mesa packages provide hardware-accelerated drivers for many popular

graphics chipsets, and Mesa, a 3D graphics application programming interface (API) compatible with the Open Graphics Library (OpenGL).

Bug Fix

BZ#821873

Previously, Mesa did not recognize Intel HD Graphics chipsets integrated into Intel E3-family processors. Consequently, these chipsets provided limited display resolutions and their graphics performance was low. This update adds support for these chipsets. As a result, the chipsets are recognized by Mesa and perform as expected.

Enhancement

BZ#788166, BZ#788168

This update adds support for AMD FirePro M100 (alternatively referred to as AMD FirePro M2000), AMD Radeon HD 74xx Series, AMD Radeon HD 75xx Series, and AMD Radeon HD 76xx Series graphics cards, and the AMD FusionA integrated graphics processing unit.

All users of xorg-x11-drv-ati and Mesa are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

5.363. XORG-X11-DRV-INTEL

5.363.1. RHBA-2012:0995 – xorg-x11-drv-intel bug fix and enhancement update

Updated xorg-x11-drv-intel packages that fix two bugs and add one enhancement are now available for Red Hat Enterprise Linux 6.

The xorg-x11-drv-intel packages contain an Intel integrated graphics video driver for the X.Org implementation of the X Window System.

Bug Fixes

BZ#692776

On Lenovo ThinkPad T500 laptops, the display could have stayed blank after opening the lid when it was used with an external display in mirror mode. Consequently, the following message appeared:

```
Could not switch the monitor configuration
Could not set the configuration for CRT63
```

With this update, the underlying source code has been modified so that the display turns on as expected when the lid is open.

BZ#711452

On Lenovo ThinkPad series laptops, the system did not always resume from the suspend state. This was dependent on monitor configuration and could occur under various circumstances, for example if the laptop was suspended docked with only external display enabled, and later resumed undocked with no external display. With this update, the system now resumes correctly regardless of the monitor configuration.

Enhancement

BZ#821521

In addition, this update adds accelerated rendering support for the Intel Core i5 and i7 processors.

All users of xorg-x11-drv-intel are advised to upgrade to these updated packages, which fix these bugs and add this enhancement.

5.364. XORG-X11-DRV-MGA

5.364.1. RHEA-2012:0940 — xorg-x11-drv-mga enhancement update

Updated xorg-x11-drv-mga packages that add an enhancement 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.

Enhancement

BZ#657580

RandR 1.2 support for G200-based graphics chipsets has been added. It allows dynamic reconfiguration of display settings to match the currently plugged in monitor. This is particularly important on servers, as they often start with no monitor attached, having it attached later in runtime.

All users of xorg-x11-drv-mga are advised to upgrade to these updated packages, which add this enhancement.

5.365. XORG-X11-DRV-QXL

5.365.1. RHSA-2013:0218 — Moderate: xorg-x11-drv-qxl security update

An updated xorg-x11-drv-qxl package that fixes one security issue is now available for Red Hat Enterprise Linux 6.

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

The xorg-x11-drv-qxl package provides 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.

Security Fix

CVE-2013-0241

A flaw was found in the way the host's qemu-kvm qxl driver and the guest's X.Org qxl driver interacted when a SPICE connection terminated. A user able to initiate a SPICE connection to a guest could use this flaw to make the guest temporarily unavailable or, potentially (if the sysctl kernel.softlockup_panic variable was set to "1" in the guest), crash the guest.

All users of xorg-x11-drv-qxl are advised to upgrade to this updated package, which contains a backported patch to correct this issue. All running X.Org server instances using the qxl driver must be restarted for this update to take effect.

5.366. XORG-X11-DRV-WACOM

5.366.1. RHBA-2012:0801 – xorg-x11-drv-wacom bug fix and enhancement update

Updated xorg-x11-drv-wacom packages that fix multiple bugs and add various enhancements 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.

The xorg-x11-drv-wacom package has been upgraded to upstream version 0.13.0, which provides a number of bug fixes and enhancements over the previous version. (BZ#752642)

Bug Fixes

BZ#734256

Prior to this update, the xorg-x11-drv-wacom driver allowed users only to use a pointer-focusing model. As a consequence, a dual-monitor layout on certain hardware could lead to an offset between the pen position and the cursor position. This update modifies the mapping offset in screen mode to provide new multi-screen handling.

BZ#802385

Prior to this update, xorg-x11-drv-wacom driver could, under certain circumstances, encounter an "off by one" error in the array access of files and a null dereference. This update modifies the array indexing and checks for the right allocation before dereferencing.

Enhancements

BZ#801319

This update adds xorg-x11-drv-wacom to HPC Compute Node (v. 6). Now, xorg-x11-drv-wacom is a dependency for the gnome-settings-daemon and the control-center.

BZ#818038

This update adds support for the Wacom Intuos4 Wireless device.

All users of xorg-x11-drv-wacom are advised to upgrade to these updated packages, which fix these bugs and add these enhancements.

5.367. XORG-X11-SERVER

5.367.1. RHSA-2012:0939 – Low: xorg-x11-server security and bug fix update

Updated xorg-x11-server packages that fix two 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 low security impact. A Common Vulnerability Scoring System (CVSS) base score, which gives a detailed severity rating, is available for each vulnerability from the CVE link(s) 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 Fixes

CVE-2011-4028

A flaw was found in the way the X.Org server handled lock files. A local user with access to the system console could use this flaw to determine the existence of a file in a directory not accessible to the user, via a symbolic link attack.

CVE-2011-4029

A race condition was found in the way the X.Org server managed temporary lock files. A local attacker could use this flaw to perform a symbolic link attack, allowing them to make an arbitrary file world readable, leading to the disclosure of sensitive information.

Red Hat would like to thank the researcher with the nickname vladz for reporting these issues.

Bug Fixes

BZ#651934, BZ#722860

Prior to this update, the KDE Display Manager (KDM) could pass invalid 24bpp pixmap formats to the X server. As a consequence, the X server could unexpectedly abort. This update modifies the underlying code to pass the correct formats.

BZ#732467

Prior to this update, absolute input devices, like the stylus of a graphic tablet, could become unresponsive in the right-most or bottom-most screen if the X server was configured as a multi-screen setup through multiple "Device" sections in the xorg.conf file. This update changes the screen crossing behavior so that absolute devices are always mapped across all screens.

BZ#748704

Prior to this update, the misleading message "Session active, not inhibited, screen idle. If you see this test, your display server is broken and you should notify your distributor." could be displayed after resuming the system or re-enabling the display, and included a URL to an external web page. This update removes this message.

BZ#757792

Prior to this update, the erroneous input handling code of the Xephyr server disabled screens on a screen crossing event. The focus was only on the screen where the mouse was located and only this screen was updated when the Xephyr nested X server was configured in a multi-screen setup. This update removes this code and Xephyr now correctly updates screens in multi-screen setups.

BZ#805377

Prior to this update, raw events did not contain relative axis values. As a consequence, clients which relied on relative values for functioning did not behave as expected. This update sets the values to the original driver values instead of the already transformed values. Now, raw events contain relative axis values as expected.

All users of xorg-x11-server are advised to upgrade to these updated packages, which correct these issues. All running X.Org server instances must be restarted for this update to take effect.

5.368. XULRUNNER

5.368.1. RHSA-2012:1361 – Critical: xulrunner security update

Updated xulrunner packages that fix one security issue are now available for Red Hat Enterprise Linux 5 and 6.

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

XULRunner provides the XUL Runtime environment for applications using the Gecko layout engine.

Security Fix

CVE-2012-4193

A flaw was found in the way XULRunner handled security wrappers. A web page containing malicious content could possibly cause an application linked against XULRunner (such as Mozilla Firefox) to execute arbitrary code with the privileges of the user running the application.

For technical details regarding this flaw, refer to the Mozilla security advisories:

http://www.mozilla.org/security/known-vulnerabilities/firefoxESR.html

Red Hat would like to thank the Mozilla project for reporting this issue. Upstream acknowledges moz_bug_r_a4 as the original reporter.

All XULRunner users should upgrade to these updated packages, which correct this issue. After installing the update, applications using XULRunner must be restarted for the changes to take effect.

5.369. YABOOT

5.369.1. RHBA-2012:0791 – yaboot bug fix update

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

The yaboot package provides a boot loader for Open Firmware based PowerPC systems. It can be used to boot IBM eServer System p machines.

Bug Fixes

BZ#711001

Prior to this update, the /etc/yaboot.conf parser failed during install when the quoted string was too long. This update modifies the code to significantly extend the size of quoted strings. Now, the /etc/yaboot.conf file parses as expected.

BZ#750199

Prior to this update, yaboot could only be build using the mock tool which creates a 32-bit PowerPC environment in chroot. This updated package supports building with rpmbuild.

All users of yaboot are advised to upgrade to this updated package, which fixes these bugs.

5.370. YUM

5.370.1. RHBA-2012:0857 – yum bug fix and enhancement update

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

Yum is a command-line utility that allows the user to check for, and automatically download and install updated RPM packages.

Bug Fixes

BZ#742363

The anacron scheduler starts the yum-cron utility with the default niceness value of 10. Consequently, Yum RPM transactions ran with a very low priority. Also, any updated service inherited this niceness value. This update adds the "reset_nice" configuration option, which allows Yum to reset the niceness value to 0 before running an RPM transaction. With this option set, Yum RPM transactions run and updated services are restarted with niceness value 0 as expected.

BZ#735234

When dependency resolving fails, yum performs RPMDB check to detect and report existing RPMDB problems. Previously, yum terminated unexpectedly if a PackageSackError exception was raised. The application now returns the message "Yum checks failed" when a PackageSackError is raised and the remaining RPMDB checks are skipped.

BZ#809392

The yum history rollback command could return a traceback if a history checksum was used for the rollback. This happened due to incorrect handling of keyword arguments in the _conv_pkg_state() function. The history checksum argument is now handled correctly.

BZ#711358

When yum was started in a directory that no longer existed, it terminated with a traceback. The yum utility now checks if the current working directory exists; if this is not the case, it changes to the root directory, and continues its execution as expected.

BZ#804120

If the "yum upgrade" command was run with the --sec-severity option arguments, the command execution could enter an infinite loop. The code has been fixed and the option works as expected.

BZ#770117

If user names and passwords for yum proxy server contained any of the characters "@", ":", or "%", they were not properly quoted in the proxy server URL and the values were misinterpreted by the HTTP client. As a result, yum failed to connect to the proxy server. This update adds proper quoting, and user names and passwords containing the characters are now resolved correctly.

BZ#809373

The Yum transactions in yum history were ordered according to their transaction time. However, this could be misleading. The transactions are now ordered according to their IDs.

BZ#769864

The "yum makecache" command could fail if one of the repositories had the "skip_if_unavailable=1" setting and was unavailable. Such repositories are now skipped as expected.

BZ#798215

The Yum utility could terminate unexpectedly with a traceback similar to the following:

```
_init__.py:2000:downloadPkgs:UnicodeEncodeError: 'ascii' codec can't encode character
```

This happened because Yum failed to handle localized error messages with UTF-8 characters generated during package downloads. UTF-8 characters in error messages are now handled correctly and localized error messages are displayed as expected.

BZ#735333

On failure, the "yum clean" command returned an incorrect error code and output containing messages that implied that yum performed the clean action successfully. The yum utility now returns only the error message and the correct error code.

BZ#817491

If the "yum provides" command was invoked with an empty-string argument, yum terminated with a traceback. The command now returns an error message and command usage information.

Enhancements

BZ#737826

Yum now prints "Verifying" messages after finishing updates, which inform the user that the respective packages were installed correctly.

BZ#690904

When run as a non-root user, yum cannot read local SSL certificate files and the download process can fail. The yum utility now checks if it can access repository certificate files. If the check fails, it returns more accurate messages containing the filename that failed the check and information that the repository was skipped.

Users of yum should upgrade to these updated packages, which fix these bugs and add these enhancements.

5.371. YUM-UTILS

5.371.1. RHBA-2012:0858 – yum-utils bug fix and enhancement update

Updated yum-utils packages that fix four bugs and add three enhancements are now available for Red Hat Enterprise Linux 6.

The yum-utils packages provide a collection of utilities and examples for the Yum package manager.

Bug Fixes

BZ#701096

Prior to this update, The reposync utility wrongly set the exit code "0" if a package was not downloaded. This update modifies the underlying code so that reposync now sets the exit code "1" if a package is either not correctly signed or fails to download.

BZ#711767

Prior to this update, the yumdownloader tool tried to download a package from all repositores that provided that particular package. As a result, after the first download a message was displayed that the file already existed. This update modifies the yumdownloader so that duplicated download attempts are now avoided.

BZ#737597

Prior to this update, the yum-debug-restore tool recognized only that the latest version of a package was installed. As a consequence, older kernel packages were not restored. This update adds support for "installonly" packages, so the whole set of installed kernel packages is restored.

BZ#782338

Prior to this update, the man page for the package-cleanup tool did not mention the changed semantics of the "--count" option. This update modifies the man page so that the "--count" option is now correctly documented.

Enhancements

BZ#684925

Prior to this update, yum could not list the dependencies and the already installed packages in the repositories that satisfy these dependencies. This update adds the "show-changed-rco" command to give a compact description of the changes to Requires, Conflicts, and Obsoletes data from installed or old files.

BZ#710579

Prior to this update, the repodiff tool only compared packages based on their name. This update adds the "--compare-arch" option to the repodiff tool to compare also the architecture.

BZ#769775

Prior to this update, the package-cleanup tool did not correctly handle kernel-PAE and kernel-xen packages. This update adds support for kernel-PAE and kernel-xen packages.

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

5.372. ZSH

5.372.1. RHBA-2012:0937 – zsh bug fix and enhancement update

Updated zsh packages that fix one bug and 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.

Bug Fix

BZ#657300

Prior to this update, the zsh shell attempted to execute mathematical expressions in the "-n" option when running in ksh mode. As a consequence, zsh emitted errors when running a syntax only check. This update modifies the source code so that mathematical expressions are now handled like any other command when using the "-n" option.

Enhancement

BZ#612685

Prior to this update, a script whose location was listed in "$PATH" could not be run with the zsh shell. With this update, users can call a script from "$PATH" with the "-o pathscript" option to search path when zsh is invoked directly.

All users are advised to upgrade to these updated packages, which fix this bug and add this enhancement.

5.373. RHNLIB

5.373.1. RHBA-2013:1212 – rhnlib bug fix update

Updated rhnlib packages that fix one bug are now available for Red Hat Enterprise Linux 6 Extended Update Support.

The rhnlib packages contain Python libraries developed specifically for interfacing with the Red Hat Network (RHN).

Bug Fix

BZ#993086

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.

5.374. RHN-CLIENT-TOOLS

5.374.1. RHBA-2013:1384 – rhn-client-tools bug fix and enhancement update

Updated rhn-client-tools packages that fix one bug and add one enhancement are now available for Red Hat Enterprise Linux 6 Extended Update Support.

Red Hat Network Client Tools provide programs and libraries that allow systems to receive software updates from Red Hat Network (RHN).

Bug Fix

BZ#993080

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#993073

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.

APPENDIX A. REVISION HISTORY

Revision 1-6.27 **Wed Jan 22 2014** **Eliška Slobodová**
Included the missing eCryptfs Technology Preview.

Revision 1-6.24 **Wed Dec 04 2013** **Milan Navrátil**
Republished the book to include Z-Stream advisories.

Revision 1-6.23 **Tue Jul 24 2013** **Miroslav Svoboda**
Republished the book to include a new kernel advisory.

Revision 1-6.21 **Fri Apr 26 2013** **Eliška Slobodová**
Republished the book to include a known issue.

Revision 1-6.20 **Wed Apr 24 2013** **Miroslav Svoboda**
Republished Technical Notes to include a new kernel advisory.

Revision 1-6.19 **Thu Apr 11 2013** **Eliška Slobodová**
Republished Technical Notes to include a kernel known issue.

Revision 1-6.18 **Thu Feb 22 2013** **Eliška Slobodová**
Republished Technical Notes to include Extended Update Support advisories relevant to Red Hat Enterprise 6.3.

Revision 1-1.17 **Wed Feb 13 2013** **Eliška Slobodová**
Updated a description in the kexec-tool erratum.

Revision 1-1.8 **Wed Jan 24 2013** **Tomáš Čapek**
Added an admonition to a kernel description.

Revision 1-1.7 **Wed Jun 20 2012** **Martin Prpič**
Release of the Red Hat Enterprise Linux 6.3 Technical Notes.

Revision 1-0 **Tue April 24 2012** **Martin Prpič**
Initial release of the Red Hat Enterprise Linux 6.3 Beta Technical Notes.