



Red Hat Enterprise Linux 5

5.0 Release Notes

Release Notes for Red Hat Enterprise Linux 5.0

Edition 1.0

Last Updated: 2017-11-21

Red Hat Enterprise Linux 5 5.0 Release Notes

Release Notes for Red Hat Enterprise Linux 5.0

Edition 1.0

Don Domingo

Engineering Services and Operations Content Services

ddomingo@redhat.com

Red Hat Enterprise Linux Documentation

Legal Notice

Copyright © 2008, 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, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

The Release Notes provide high-level coverage of the improvements and additions that have been implemented in Red Hat Enterprise Linux 5.0. For detailed documentation on all changes to Red Hat Enterprise Linux for the 5.0 update, refer to the Technical Notes.

Table of Contents

| | |
|------------------------------------------|-----------|
| 1. RELEASE NOTES FOR X86 | 3 |
| 1.1. Installation-Related Notes | 3 |
| 1.1.1. ISO Contents and Registration | 3 |
| 1.1.2. New RPM GPG Signing Keys | 4 |
| 1.1.3. Subversion | 4 |
| 1.1.4. Other Installation Notes | 4 |
| 1.2. Technology Previews | 5 |
| 1.3. Known Issues | 8 |
| 1.4. General Information | 13 |
| 1.5. Driver Update Program | 20 |
| 1.6. Internationalization | 21 |
| 1.7. Kernel Notes | 23 |
| 2. RELEASE NOTES FOR X86_64 | 28 |
| 2.1. Installation-Related Notes | 28 |
| 2.1.1. ISO Contents and Registration | 29 |
| 2.1.2. New RPM GPG Signing Keys | 29 |
| 2.1.3. Subversion | 29 |
| 2.1.4. Other Installation Notes | 30 |
| 2.2. Technology Previews | 30 |
| 2.3. Known Issues | 34 |
| 2.4. General Information | 37 |
| 2.5. Driver Update Program | 42 |
| 2.6. Internationalization | 43 |
| 2.7. Kernel Notes | 45 |
| 3. RELEASE NOTES FOR IA64 | 50 |
| 3.1. Installation-Related Notes | 51 |
| 3.1.1. ISO Contents and Registration | 51 |
| 3.1.2. New RPM GPG Signing Keys | 51 |
| 3.1.3. Subversion | 52 |
| 3.1.4. Other Installation Notes | 52 |
| 3.2. Technology Previews | 53 |
| 3.3. Known Issues | 58 |
| 3.4. General Information | 62 |
| 3.5. Driver Update Program | 67 |
| 3.6. Internationalization | 68 |
| 3.7. Kernel Notes | 70 |
| 4. RELEASE NOTES FOR PPC | 76 |
| 4.1. Installation-Related Notes | 76 |
| 4.1.1. ISO Contents and Registration | 76 |
| 4.1.2. New RPM GPG Signing Keys | 77 |
| 4.1.3. Subversion | 77 |
| 4.1.4. Other Installation Notes | 78 |
| 4.2. Technology Previews | 78 |
| 4.3. Known Issues | 81 |
| 4.4. General Information | 85 |
| 4.5. Driver Update Program | 90 |
| 4.6. Internationalization | 91 |
| 4.7. Kernel Notes | 93 |

| | |
|-----------------------------------------|-----------|
| 5. RELEASE NOTES FOR S390X | 98 |
| 5.1. Installation-Related Notes | 98 |
| 5.1.1. ISO Contents and Registration | 99 |
| 5.1.2. New RPM GPG Signing Keys | 99 |
| 5.1.3. Subversion | 100 |
| 5.1.4. Other Installation Notes | 100 |
| 5.2. Technology Previews | 101 |
| 5.3. Known Issues | 104 |
| 5.4. General Information | 107 |
| 5.5. Driver Update Program | 113 |
| 5.6. Internationalization | 114 |
| 5.7. Kernel Notes | 116 |

1. RELEASE NOTES FOR X86

1.1. Installation-Related Notes

The following section includes information specific to the installation of Red Hat Enterprise Linux and the Anaconda installation program.



NOTE

In order to upgrade an already-installed Red Hat Enterprise Linux, you must use Red Hat Network to update those packages that have changed.

You may use Anaconda to perform a fresh installation of Red Hat Enterprise Linux 5 or to perform an upgrade from *the latest updated version* of Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5.

If you are copying the contents of the Red Hat Enterprise Linux 5 CD-ROMs (in preparation for a network-based installation, for example) be sure to copy the CD-ROMs for the operating system *only*. Do not copy the Supplementary CD-ROM, or any of the layered product CD-ROMs, as this will overwrite files necessary for Anaconda's proper operation. These CD-ROMs must be installed *after* Red Hat Enterprise Linux has been installed.

1.1.1. ISO Contents and Registration

The organization of software component packages into product-specific variants has changed from previous versions of Red Hat Enterprise Linux. The total number of different variants and ISO images has been reduced to two:

- Red Hat Enterprise Linux 5 Server
- Red Hat Enterprise Linux 5 Client

The ISO images contain software packages for a number of optional repositories that provide additional functionality over the core distribution, such as Virtualization, Clustering or Cluster Storage. For more information about the Server variants, Client variants and available options, please refer to <http://www.redhat.com/rhel/>.

With optional content in the same tree or ISO image, it is important to avoid a mismatch between the components offered for installation and those covered by the subscription. Such a mismatch could result in an increased exposure to bug and vulnerability risks.

In order to ensure that the components offered are in sync with the subscription, Red Hat Enterprise Linux 5 requires entering an *Installation Number* that will be used to configure the installer to offer the right package set. This Installation Number is included in your subscription.

If you skip entering the Installation Number, this will result in a core Server or Desktop installation. Additional functionality can then be added manually at a later time. For more information about Installation Numbers, please refer to <http://www.redhat.com/apps/support/in.html>.

The Installation Number used during the installation process will be saved in `/etc/sysconfig/rhn/install-num`. When registering with Red Hat Network, this file will be

referenced by `rhn_register` to automatically determine which appropriate child channels the system should be subscribed to.

1.1.2. New RPM GPG Signing Keys

A new release signing key is used to sign Red Hat Enterprise Linux 5 packages. When updating a system for the first time, you will be prompted to allow this key to be installed.

Signing keys are distributed in the following files:

- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release` – contains the public key for the new release signing key
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-auxiliary` – contains the public key for an auxiliary release signing key, currently not in use
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-former` – contains the public key for the previous release signing key, used for past Red Hat Enterprise Linux releases

1.1.3. Subversion

In Red Hat Enterprise Linux 5, the Subversion version control system is linked against Berkeley DB 4.3. If you are upgrading from Red Hat Enterprise Linux 4 and any Subversion repositories which use the Berkeley DB backend "BDB" (rather than the pure file system-based "FSFS" backend) have been created on the system, special care must be taken to ensure the repositories can be accessible after the upgrade. The following process must be performed on the Red Hat Enterprise Linux 4 system, *prior* to upgrading to Red Hat Enterprise Linux 5:

1. Shut down any running processes and ensure that no processes can access the repository (for example, `ht tpd`, `svnserve` or any local users with direct access).
2. Create a backup of the repository using the following command:

```
svnadmin dump /path/to/repository | gzip > repository-backup.gz
```

3. Run the `svnadmin recover` command on the repository:

```
svnadmin recover /path/to/repository
```

4. Delete any unused log files in the repository:

```
svnadmin list-unused-dblogs /path/to/repository | xargs rm -vf
```

5. Delete any remaining shared-memory files in the repository:

```
rm -f /path/to/repository/db/___db.0*
```

1.1.4. Other Installation Notes

- When installing Red Hat Enterprise Linux 5 using split installation media (for example, CD or NFSISO), an error occurs during the installation of `amanda-server`.

As such, if you wish to use `amanda-server`, you should install it using `yum` after Red Hat Enterprise Linux 5 is installed.

Note that this issue does not affect installations where non-split media is used (for example, DVD or NFS tree installations).

- If IDE/PATA (*Parallel ATA*) devices are configured in "100% Native" mode, some BIOSes may prevent the Red Hat Enterprise Linux 5 installation process from completing successfully. To prevent this from occurring, configure the IDE/PATA mode as "Legacy" in the BIOS.
- The *IBM System z* does not provide a traditional Unix-style physical console. As such, Red Hat Enterprise Linux 5 for the *IBM System z* does not support the `firstboot` functionality during initial program load.

To properly initialize setup for Red Hat Enterprise Linux 5 on the *IBM System z*, run the following commands after installation:

- `/usr/bin/setup` – provided by the `setuptools` package
- `/usr/bin/rhn_register` – provided by the `rhn-setup` package
- During the installation process, Anaconda will automatically choose the kernel package to be installed. The kernel selected by default does not allow Red Hat Enterprise Linux 5 to detect more than 4GB of RAM. As such, if your system has more than 4GB of RAM, you need to install the kernel-PAE variant of the kernel after installation.

Note that this does not apply when a virtual install is performed.

- When booting Anaconda with PXE using the parameter `ksdevice=bootif`, you will still be prompted for the ethernet interface to use during installation. If only one ethernet device is plugged in, use the parameter `ksdevice=link` instead. Alternatively, you can also specify the interface manually.

1.2. Technology Previews

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 5 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. Erratas 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.

Stateless Linux

Included in this release of Red Hat Enterprise Linux 5 are enabling infrastructure pieces for *Stateless Linux*. Stateless Linux is a new way of thinking about how a system is to be run and managed, designed to simplify provisioning and management of large numbers of systems by making them easily replaceable. This is accomplished primarily by establishing prepared system images which get replicated and managed across a large number of stateless systems, running the operating system in a read-only manner (please refer to `/etc/sysconfig/readonly-root` for more details).

In its current state of development, the Stateless features are subsets of the intended goals. As such, the capability is being labeled as a technology preview.

The following is a list of the initial capabilities included in Red Hat Enterprise Linux 5:

- running a stateless image over NFS
- running a stateless image via loopback over NFS
- running on iSCSI

It is highly recommended that those interested in testing stateless code read the HOWTO at <http://fedoraproject.org/wiki/StatelessLinuxHOWTO> and join stateless-list@redhat.com.

GFS2

GFS2 is an evolutionary advancement based on the GFS file system. While fully functional, GFS2 is not yet considered production-ready. GFS2 is targeted to become fully supported in a subsequent Red Hat Enterprise Linux 5 update. There is also an in-place conversion utility, `gfs2_convert`, which can update the metadata of the older GFS file system format, converting it to a GFS2 file system.

FS-Cache

FS-Cache is a local caching facility for remote file systems that allows users to cache NFS data on a locally mounted disk. To set up the FS-Cache facility, install the `cachefilesd` RPM and refer to the instructions in `/usr/share/doc/cachefilesd-<version>/README`.

Replace `<version>` with the corresponding version of the `cachefilesd` package installed.

Compiz

Compiz is an OpenGL-based compositing window manager. In addition to regular window management, `compiz` also acts as a compositing manager, coordinating and synchronizing the overall desktop redrawing to provide a smoother desktop experience with less flicker.

`Compiz` uses 3D hardware acceleration to render effects such as live thumbnail windows, window drop shadows, animated window minimizing and transitions between virtual desktops.

Due to limitations in the current rendering architecture, `compiz` cannot work correctly with direct rendering OpenGL applications or applications using the Xv extension. Such applications will exhibit harmless rendering artifacts; because of this, `compiz` is currently a technology preview.

Enhancement for Ext3

In Red Hat Enterprise Linux 5, the EXT3 file system capacity has been extended beyond 8TB to a maximum of 16TB. This capability is being included as a technology preview, and is targeted for full support in a future release of Red Hat Enterprise Linux 5.

AIGLX

AIGLX is a technology preview feature of the otherwise fully supported X server. It aims to enable GL-accelerated effects on a standard desktop. The project consists of the following:

- a lightly modified X server
- an updated Mesa package that adds new protocol support

By installing these components, you can have GL-accelerated effects on your desktop with very few

changes, as well as the ability to enable and disable them at will without replacing your X server. AIGLX also enables remote GLX applications to take advantage of hardware GLX acceleration.

Frysk GUI

The goal of the *frysk* project is to create an intelligent, distributed, always-on system monitoring and debugging tool that allows developers and system administrators to:

- monitor running processes and threads (including creation and destruction events)
- monitor the use of locking primitives
- expose deadlocks
- gather data
- debug any given process by choosing it from a list or allowing *frysk* to open a source code (or other) window on a process that is crashing or misbehaving

In Red Hat Enterprise Linux 5 the *frysk* graphical user interface is a technology preview, whereas the *frysk* command line interface is fully supported.

Systemtap

Systemtap provides free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. With the help of *systemtap*, developers no longer need to go through the tedious and disruptive instrument, recompile, install, and reboot sequence that may be otherwise required to collect data.

Dogtail

Dogtail is a GUI test tool and automation framework written in Python that uses *Accessibility* technologies to communicate with desktop applications.

Support for Indic Languages and Sinhalese

Red Hat Enterprise Linux 5 also features support for the following languages as technology preview:

- *Assamese*
- *Kannada*
- *Sinhalese*
- *Telugu*

For more information about how to install and enable support for these languages, refer to the *Internationalization* section of this document.

Installing to dm-multipath Devices

Anaconda now has the capability to detect, create, and install to *dm-multipath* devices. To enable this feature, add the parameter `mpath` to the kernel boot line.

Note that the parameter `mpath` may cause a boot failure if a device's `major:minor` number changes. This issue will be addressed in a future update of Red Hat Enterprise Linux 5.

Installation / Boot for iSCSI software initiator (open-iscsi)

Anaconda now provides the ability to install to an iSCSI device. Booting and installing is fully supported with the *QLogic qla4xxx* hardware initiator. However, the capability to install to an iSCSI device for the `open-iscsi` software initiator is currently considered a Technology Preview, due to the following issues:

- Text mode installation does not complete. You must do a graphical install, or an automated kickstart install.
- Media-based installations do not complete. You must do a network-based install.
- Depending on the timing of events, Anaconda may be unable to detect all the iSCSI targets or LUNs. When this occurs, use the installer shell to configure the storage through iSCSI commands.
- The `iscsid` daemon may not properly start. Such an occurrence will prevent the system from handling all iSCSI errors, such as network problems, SCSI/iSCSI timeouts, and target errors. To confirm that the `iscsid` daemon is running, run the command `iscsiadm -m session -i` and check that the line, **Internal iscsid Session State:** reports a value (it can be any value).
- On certain iSCSI target implementations, the system may hang during shutdown.
- On certain iSCSI target implementations, the system may hang during reboot. To avoid this, shutdown the system and boot it up again (instead of rebooting directly from a session).
- Booting from iSCSI devices on the *IBM System p* does not work reliably. While installation on an iSCSI device may appear to succeed, the resulting installation will not boot properly.
- On the first boot after install, you may receive SELinux errors such as the following:

```
kernel: audit(1169664832.270:4): avc: denied { read
} for pid=1964 comm="iscsid"
```

To work around this, boot the system with the kernel parameter `enforcing=0`. Once the system has properly booted, use the command `setenforce 1` to restore enforcing mode.

These limitations will be addressed in a future Red Hat Enterprise Linux 5 update.

1.3. Known Issues

- Host bus adapters that use the MegaRAID driver must be set to operate in "Mass Storage" emulation mode, not in "I2O" emulation mode. To do this, perform the following steps:
 1. Enter the MegaRAID **BIOS Set Up Utility**.
 2. Enter the **Adapter settings** menu.
 3. Under **Other Adapter Options**, select **Emulation** and set it to **Mass Storage**.

If the adapter is incorrectly set to "I2O" emulation, the system will attempt to load the `i2o` driver. This will fail, and prevent the proper driver from being loaded.

Previous Red Hat Enterprise Linux releases generally do not attempt to load the I2O driver before the MegaRAID driver. Regardless of this, the hardware should always be set to "Mass Storage" emulation mode when used with Linux.

- When you install a fully virtualized guest configured with `vcpus=2`, the fully virtualized guest may take an unreasonably long time to boot up.

To work around this, destroy the slow-booting guest using the command `xm destroy <guest id>` and then use `xm create <guest id>` to start the same guest afterwards.

- Red Hat Enterprise Linux 5 includes `openmpi-1.1.1-4.el5` (from the OFED 1.1 distribution), which has been discovered to eventually quit working entirely. This happens after the `openmpi` stack works as expected for a varying amount of time.

For updated versions of `openmpi`, please check <http://people.redhat.com/dledford/Infiniband/openmpi>

- Installing Windows Server 2003 as a guest on a fully virtualized Red Hat Enterprise Linux 5 system ends unexpectedly after completing the first stage of installation. When this happens, the graphical console window closes, and the guest disappears from the **Virtual Machine Manager's** list of machines, resulting in a **Broken pipe** error.

This issue will be resolved in an upcoming Red Hat Enterprise Linux 5 update. To work around this, use the following command at the terminal:

```
xm create /etc/xen/<name of guest machine>
```

Afterwards, open the virtual machine.

- When attempting to create a fully virtualized Windows Server 2003 from a CD / DVD, the second stage of the guest install will not continue upon reboot.

To work around this, edit `/etc/xen/<name of guest machine>` by properly appending an entry for the CD / DVD device.

If an installation to a simple file is used as a virtual device, the `disk` line of `/etc/xen/<name of guest machine>` will read like the following:

```
disk = [ 'file:/PATH-OF-SIMPLE-FILE,hda,w' ]
```

A DVD-ROM device located on the host as `/dev/dvd` can be made available to stage 2 of the installation as `hdc` by appending an entry like `'phy:/dev/dvd,hdc:cdrom,r'`. As such, the disk line should now read as follows:

```
disk = [ 'file:/opt/win2003-sp1-20061107,hda,w',
        'phy:/dev/dvd,hdc:cdrom,r' ]
```

The precise device path to use may vary depending on your hardware.

- `rmmod xennet` causes `domU` to crash; this is caused by a grant table issue in the Virtualization feature. Due to the current inability of the Virtualization feature to asynchronously release grant table operations, it is unsafe to unload the `xennet` module in guests. In such situations, grant tables are used to perform backend-frontend communication, and there is no guarantee that the backend will release the references, leading to an inevitable memory leak.

This issue will be resolved in the next minor release of Red Hat Enterprise Linux 5. At present, users are advised *not* to unload the `xennet` module in guests.

- Running `ethtool eth0` outputs incomplete information about the ethernet card settings. This only occurs in systems running a virtualized kernel, since the Virtualization feature uses a networking setup where the physical ethernet device is identified as `peth0`. As such, the correct command for retrieving information about the physical ethernet device is `ethtool peth0`.
- Installing the Virtualization feature may cause a `time went backwards` warning on HP systems with model numbers `xw9300` and `xw9400`.

To work around this issue for `xw9400` machines, configure the BIOS settings to enable the **HPET** timer. Note that this option is not available on `xw9300` machines.

HP will notify `xw9300` and `xw9400` users when a new BIOS image is available.

- When using Red Hat Enterprise Linux 5 on a machine with an *nVidia CK804* chipset installed, you may receive kernel messages similar to the following:

```
kernel: assign_interrupt_mode Found MSI capability
kernel: pcie_portdrv_probe->Dev[005d:10de] has invalid IRQ. Check
vendor BIOS
```

These messages indicate that certain PCI-E ports are not requesting IRQs. Further, these messages do not, in any way, affect the operation of the machine.

- Some *Cisco Aironet* Wireless devices prevent **NetworkManager** from storing connection details for wireless networks that do not broadcast an SSID. This is caused by a *Cisco Aironet* Wireless device firmware limitation.
- Laptops that have the *Cisco Aironet MPI-350* wireless card equipped may hang trying to get a DHCP address during any network-based installation using the wired ethernet port.

To work around this, use local media for your installation. Alternatively, you can disable the wireless card in the laptop BIOS prior to installation (you can re-enable the wireless card after completing the installation).

- Currently, `system-config-kickstart` does not support package selection and deselection. When using `system-config-kickstart`, the **Package Selection** option indicates that it is disabled. This is because `system-config-kickstart` uses `yum` to gather group information, but is unable to configure `yum` to connect to Red Hat Network.

This issue is currently being investigated for resolution by the next minor release of Red Hat Enterprise Linux 5. At present, you need to update package sections in your kickstart files manually. When using `system-config-kickstart` to open a kickstart file, it will preserve all package information in it and write it back out when you save.

- Systems with SATA controllers may pause during the boot process, displaying the following error message:

```
ata2: port is slow to respond, please be patient
```

Afterwards, the following error message appears:

```
ata2: reset failed, giving up
```

Note that after the second error message, the system will continue the normal boot process. Other than the delay, there is no impact to the system; as long as the SATA drives are physically present they will still be detected properly.

- 4-socket AMD *Sun Blade X8400 Server Modules* systems that do not have memory configured in `node 0` will panic during boot. Systems should be configured with memory in `node 0` to prevent the kernel panic.
- Installing to LVM mirror devices through Anaconda is currently not supported. This capability will be added in a future update of Red Hat Enterprise Linux 5.
- When installing Red Hat Enterprise Linux 5 from a directory on an NFS server containing Red Hat Enterprise Linux ISO images, Anaconda may display the following error message:

```
Unable to read package metadata. This may be due to a missing
repopdata directory.
Please ensure that your install tree has been correctly generated.
Cannot
open/read repomd.xml file for repository:
```

This problem occurs if the directory holding the ISO images also contains a partially unpacked installation tree (for example, the `/images` directory from the first ISO). The presence of such directories results in the error stated above.

To prevent this error, unpack trees only to directories other than the one containing the installation ISO images.

- Boot-time logging to `/var/log/boot.log` is not available in this release of Red Hat Enterprise Linux 5. An equivalent functionality will be added in a future update of Red Hat Enterprise Linux 5.
- Neither `kexec` nor `kdump` are able to dump onto disks attached to an `accraid` controller.

To work around this issue, use `scp` for network dumping. Alternatively, you can also dump onto a disk through a different controller.

- Running `tvtime` and `xawtv` with the `bttv` kernel module causes the system to freeze. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.

To work around this, add the parameter `mem=3000m` to the kernel boot line.

- When using an *IBM T43* laptop on a docking station attached to a secondary Dell monitor, both laptop and secondary screen display incorrect resolutions upon boot-up.

To prevent this from occurring in future logins, perform the following steps:

1. Open the **Display settings** menu using the command `system-config-display`.
2. Click the **Dual head** tab.
3. Check **Use dual head** and enter the proper configuration for the secondary monitor.
4. Reboot the system.

- The Supplementary CD of this release contains the **Mozilla** plugins **flash-plugin** and **acroread-plugin**. Both of these plug-ins are 32-bit, and as such it is recommended that they not be installed with the 64-bit **Firefox** browser.
- Installing a fully virtualized guest using split installation media -- specifically, multiple CD-ROMs -- may fail when required to switch between installation CDs. During the guest OS installation process, users may be prevented from mounting or ejecting installation CDs, which prevents the installation from completing.

As such, it is recommended that you use the QEMU monitor console to switch CD-ROM images during the guest OS installation process. The procedure is as follows:

1. Open a graphical VNC console to the guest OS.
2. Unmount the CD-ROM device in the guest OS.
3. Switch to the QEMU monitor console by pressing **Ctrl-Alt-2**.
4. Run the command `eject hdc`.
5. Run the command `change hdc <path to the CD-ROM in host system>`.
6. Switch back to the guest OS console by pressing **Ctrl-Alt-1**.
7. Mount the CD-ROM device in the guest OS.

Note that when using a regular VNC client the host X server may encounter some difficulty interpreting the **Ctrl-Alt-2** and **Ctrl-Alt-1** command. To work around this in **virt-manager**, use *sticky keys*. Pressing **Ctrl** three times makes it "sticky" until the next non-modifier is pressed. As such, to send **Ctrl-Alt-1**, press **Ctrl** twice before pressing **Ctrl-Alt-1**.

- Some machines that use **NVIDIA** graphics cards may display corrupted graphics or fonts when using the graphical installer or during a graphical login. To work around this, switch to a virtual console and back to the original X host.
- The Red Hat Enterprise Linux 5 *Driver Update Model* creates modified **initrd** images whenever a **kmod** package that includes a bootpath-modifying driver is installed. In time, the number of backup **initrd** images may soon fill the **/boot** partition, particularly if the system undergoes a sizeable number of driver updates.

As such, it is recommended that you monitor the free space on the **/boot** partition if you regularly perform driver updates. You can free up more space in **/boot** by removing older **initrd** images; these files end in **.img0**, **.img1**, **.img2**, and so on.

- The Red Hat Enterprise Linux virtualization kernel may not work correctly with more than 64GB of memory. To boot the virtualization kernel on machines that have more than 64GB of physical memory installed, you may need to add **dom0_mem=4G mem=64G** to the kernel command-line. For example:

```
title Red Hat Enterprise Linux Server (2.6.18-4.el5xen)
  root (hd0,0)
  kernel /xen.gz-2.6.18-4.el5 dom0_mem=4G mem=64G
  module /vmlinuz-2.6.18-4.el5xen ro root=LABEL=/
  module /initrd-2.6.18-4.el5xen.img
```


- Autorun on removable media is currently disabled. To install packages from the Red Hat Enterprise Linux Supplementary CD, launch the CD installer manually using the following command:

```
system-cdinstall-helper /media/path-to-mounted-drive
```

- When upgrading from Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5, the Deployment Guide is not automatically installed. You need to use `pirut` to manually install it after completing the upgrade.
- An `autofs` bug prevents multi-mounts from working properly.

During an expiry, if the last multi-mount component to be checked does not have a mount associated with it while other components are busy, `autofs` erroneously determines the multi-mount to be expirable. This causes the multi-mount to be partially expired, resulting in the multi-mount becoming unresponsive to further mount requests and expire runs.

To permanently resolve this problem, update `autofs` using the command `yum update autofs`.

- The system may not successfully reboot into a `kexec/kdump` kernel if X is running and using a driver other than `vesa`. This problem only exists with *ATI Rage XL* graphics chipsets.

If X is running on a system equipped with *ATI Rage XL*, ensure that it is using the `vesa` driver in order to successfully reboot into a `kexec/kdump` kernel.

- Creating a fully virtualized guest using a `boot.iso` on an NFS share mounted as read-write will not complete correctly. To work around this problem, mount the NFS share as read-only.

If you are unable to mount the NFS share as read-only, copy the `boot.iso` to the local `/var/lib/xen/images/` directory.

1.4. General Information

This section contains general information not specific to any other section of this document.

Red Hat Enterprise Linux Deployment Guide

This release of Red Hat Enterprise Linux includes a fairly comprehensive Deployment Guide. To access it, go to **System** (on the top panel) => **Documentation** => **Red Hat Enterprise Linux Deployment Guide**.

It is the intention of Red Hat to provide fully localized versions of the Deployment Guide for all supported languages. If you have installed a localized version of the Deployment Guide, it is recommended that you update it when a new version becomes available through Red Hat Network.

Virtualization

Red Hat Enterprise Linux 5 features Xen-based virtualization capabilities for i686 and x86-64, as well as the software infrastructure needed to manage a virtualized environment.

The implementation of Virtualization in Red Hat Enterprise Linux 5 is based on the *hypervisor*, which facilitates extremely low overhead virtualization through paravirtualization. With Intel Virtualization Technology or AMD AMD-V capable processors, virtualization in Red Hat Enterprise Linux 5 also allows operating systems to run unmodified in fully virtualized mode.

Virtualization on Red Hat Enterprise Linux 5 also features the following:

- *Libvirt*, a library that provides a consistent, portable API for managing virtual machines.
- *Virtual Machine Manager*, a graphical utility for monitoring and managing virtual machines.
- Virtual machine support in the installer, including the ability to kickstart virtual machines.

Red Hat Network also supports virtual machines.

At present, the Virtualization feature has the following limitations:

- When Virtualization is enabled, neither suspend to RAM nor suspend to disk are supported, and CPU frequency scaling cannot be performed.
- Hardware-virtualized guests cannot have more than 2GB of virtual memory.
- Hardware-virtualized guests cannot have more than 4GB of virtual memory.
- Fully virtualized guests cannot be saved, restored or migrated.
- The `xm create` command does not have a graphical equivalent in Virtual Machine Manager.
- Virtualization only supports the *bridged networking* component. All corresponding tools used by guests automatically choose this as the default.
- The default Red Hat SELinux policy for Virtualization only allows configuration files to be written to `/etc/xen`, log files to be written to `/var/log/xen/`, and disk files (including core dumps) to be written to `/var/lib/xen`. These defaults can be changed using the `semanage` tool.
- The hypervisor included in this release of Virtualization is *not* NUMA-aware; as such, its performance on NUMA machines may be sub-optimal. This will be addressed in a future update of Red Hat Enterprise Linux 5.

To work around this, enable `memory node interleaving` in the NUMA machine's BIOS. This ensures a more consistent performance.

- Paravirtualized domains currently do not support keymaps other than `en-US`. As such, other keyboards may not be able to type certain keystrokes. This will be addressed in a future update of Red Hat Enterprise Linux 5.
- The virtualized kernel cannot use the `kdump` function.
- `qcow` and `vmdk` images are not supported. When manually configuring guests, images backed by a physical or logical device should use the `phy: type`. For file-backed images, set the image type to `tap:aio:` for paravirtualized guests and `file:` for fully virtualized guests.
- Profiling of fully virtual domains may be inaccurate. This issue will be addressed in the next minor release of Red Hat Enterprise Linux 5.
- Paravirtualized domains can only auto-detect relative mouse movement, and pointer movement is rather erratic. This will be addressed in a future update of Red Hat Enterprise Linux 5.

- Some `dom0` serial console setups may require additional configuration. Refer to the *Troubleshooting* section of the *Virtualization Guide* for more details on recommended configurations.
- In order to have a working console for a paravirtualized guest, you need to specify `console=xvc0` in the kernel command line.
- When guest operating systems are configured to use sparse files, `dom0` can run out of disk space. Such occurrences prevent guest disk writes from completing, and can cause data loss in guests. Further, guests that use sparse files do not synchronize I/O safely.

As such, it is recommended that you use non-sparse files instead. To configure guests to use non-sparse files, use the option `--nonsparse` when conducting a `virt-install`.

Web Server Packaging Changes

Red Hat Enterprise Linux 5 now includes version 2.2 of the Apache HTTP Server. This release brings a number of improvements over the 2.0 series, including:

- improved caching modules (`mod_cache`, `mod_disk_cache`, `mod_mem_cache`)
- a new structure for authentication and authorization support, replacing the authentication modules provided in previous versions
- support for proxy load balancing (`mod_proxy_balancer`)
- support for handling large files (namely, greater than 2GB) on 32-bit platforms

The following changes have been made to the default `httpd` configuration:

- The `mod_cern_meta` and `mod_asis` modules are no longer loaded by default.
- The `mod_ext_filter` module is now loaded by default.

If you are upgrading from a previous release of Red Hat Enterprise Linux, the `httpd` configuration will need to be updated for `httpd 2.2`. For more information, refer to <http://httpd.apache.org/docs/2.2/upgrading.html>.

Note that any third-party modules compiled for `httpd 2.0` must be rebuilt for `httpd 2.2`.

php

Version 5.1 of PHP is now included in Red Hat Enterprise Linux 5, which includes a number of changes to the language along with significant performance improvements. Some scripts might need to be edited for use with the new version; please refer to the link below for more information on migrating from PHP 4.3 to PHP 5.1:

<http://www.php.net/manual/en/migration5.php>

The `/usr/bin/php` executable is now built using the CLI command-line SAPI, rather than the CGI SAPI. Use `/usr/bin/php-cgi` for CGI SAPI. The `php-cgi` executable also includes FastCGI support.

The following extension modules have been added:

- the `mysqli` extension, a new interface designed specifically for MySQL 4.1 (included in the `php-mysql` package)

- `date`, `hash`, `Reflection`, `SPL` and `SimpleXML` (built-in with the `php` package)
- `pdo` and `pdo_sqlite` (in the `php-pdo` package)
- `pdo_mysql` (in the `php-mysql` package)
- `pdo_pgsql` (in the `php-pgsql` package)
- `pdo_odbc` (in the `php-odbc` package)
- `soap` (in the `php-soap` package)
- `xmlreader` and `xmlwriter` (in the `php-xml` package)
- `dom` (replacing the `domxml` extension in the `php-xml` package)

The following extension modules are no longer included:

- `dbx`
- `dio`
- `yp`
- `overload`
- `domxml`

The PEAR Framework

The PEAR framework is now packaged in the `php-pear` package. Only the following PEAR components are included in Red Hat Enterprise Linux 5:

- `Archive_Tar`
- `Console_Getopt`
- `XML_RPC`

Encrypted Swap Partitions and Non-root File Systems

Red Hat Enterprise Linux 5 now provides basic support for encrypted swap partitions and non-root file systems. To use these features, add the appropriate entries to `/etc/crypttab` and reference the created devices in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry:

```
my_swap /dev/hdb1 /dev/urandom swap,cipher=aes-cbc-essiv:sha256
```

This creates the encrypted block device `/dev/mapper/my_swap`, which can be referenced in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry for a file system volume:

```
my_volume /dev/hda5 /etc/volume_key cipher=aes-cbc-essiv:sha256
```

The `/etc/volume_key` file contains a plaintext encryption key. You can also specify `none` as the key file name; this configures the system to ask for the encryption key during boot instead.

It is recommended to use LUKS (Linux Unified Key Setup) for setting up file system volumes. To do this, follow these steps:

1. Create the encrypted volume using `cryptsetup luksFormat`.
2. Add the necessary entry to `/etc/crypttab`.
3. Set up the volume manually using `cryptsetup luksOpen` (or reboot).
4. Create a file system on the encrypted volume.
5. Add the necessary entry to `/etc/fstab`.

mount and umount

The `mount` and `umount` commands no longer directly support NFS; a built-in NFS client no longer exists. A separate `nfs-utils` package, which provides `/sbin/mount.nfs` and `/sbin/umount.nfs` helpers, must be installed for this.

CUPS Printer Browsing

CUPS printer browsing over a local subnet can be configured using the graphical tool `system-config-printer`. It can also be done using the CUPS web interface, <http://localhost:631/>.

To use directed broadcasts for printer browsing between subnets, open `/etc/cups/cupsd.conf` on the clients and replace `BrowseAllow @LOCAL` with `BrowseAllow ALL`.

ATI and R500 Support

ATI graphics cards based on the R500 chipset are supported for the `vesa` driver only, and are not supported by Red Hat Enterprise Linux 5 on external monitors, LCD projectors or accelerated 3D support.

up2date and yum

`up2date` is being deprecated in favor of `yum` (Yellowdog Updater Modified). As such, it is advisable that you revise any `up2date`-dependent scripts your system is using accordingly. For more information about `yum`, consult its man page with the command `man yum`; you can also consult the installed documentation under the directories `/usr/share/doc/yum-<version>` and `/usr/share/doc/yum-metadata-parser-<version>` (replace `<version>` with the corresponding version of `yum` and `yum-metadata-parser` installed).

OpenLDAP Server and Red Hat Directory Server

Red Hat Directory Server is an LDAP-based server that centralizes enterprise and network data into an OS-independent, network-based registry. It is set to replace OpenLDAP server components, which will be deprecated *after* Red Hat Enterprise Linux 5. For more information about Red Hat Directory Server, refer to <http://www.redhat.com/software/rha/directory/>.

i810 Driver and i830 Support

The `i810` driver supports all integrated Intel graphics chipsets, from `i810` to `i965`. However, the support for `i830` (and newer) chipsets is limited; the `i810` driver can only set modes listed in the video BIOS. If your machine has an `i830` or newer chipset installed, run the following command to determine what the available modes are:

grep Mode: `/var/log/Xorg.0.log`

Modes marked with an asterisk (*) are available for selection.

Many laptop video BIOSes do not supply a mode that matches the native panel size. Therefore the chosen mode may appear stretched, distorted, or with black borders. As such, if your chosen mode does not display properly, you need a BIOS update from your hardware vendor for the native panel size to work correctly.

Smart Card Login

Red Hat Enterprise Linux 5 includes support for *Smart Cards*, which provide secure storage for your key pair and an associated public key certificate. These keys are protected through a PIN that you need to input when a key or certificate on the Smart Card is required.

Deploying Smart Cards in a Red Hat Enterprise Linux 5 environment allows you to leverage features such as Kerberos and S/MIME to increase security in relation to authentication. Red Hat Enterprise Linux 5 supports the following:

- Axalto Cyberflex 32K e-Gate
- DoD CAC Cards

To set up Smart Card authentication, your network will need to be equipped with Red Hat Directory Server and Red Hat Certificate System. For more information regarding Smart Cards, refer to the Red Hat Enterprise Linux Deployment Guide Chapter on *Single Sign-On*.

Intel PRO/Wireless 3945ABG Network Connection Support

This release of Red Hat Enterprise Linux 5 includes support for the *ipw3945 (Intel PRO/Wireless 3945ABG Network Connection)* adapter. The Red Hat Enterprise Linux 5 Supplementary disc contains the driver, regulatory daemon and firmware needed to support this adapter.

To enable support for the ipw3945 wireless adapter, search the Red Hat Enterprise Linux 5 Supplementary disc for packages with filenames containing "3945" and install them.

rawio

rawio is a deprecated interface; however, Red Hat Enterprise Linux 5 still includes support for it. If you have an application that performs device access using *rawio*, it is highly recommended that you modify your application to open the block device with the `O_DIRECT` flag. The *rawio* interface will remain throughout the life of Red Hat Enterprise Linux 5, but is a candidate for removal in a future release.

Currently, AIO (*Asynchronous I/O*) on file systems is only supported in `O_DIRECT` or non-buffered mode. Further, note that the asynchronous poll interface is no longer present, and that AIO on pipes is no longer supported.

ctmptc

ctmptc is a deprecated driver; however, it will still be included throughout the life of Red Hat Enterprise Linux 5. Note that it is a candidate for removal from future releases.

Policy Modules and semanage Support

Red Hat Enterprise Linux 5 now supports *policy modules* and *semanage*. Policy modules simplify the creation and distribution of policy customizations and third-party policies through the use of the `semodule` and `checkmodule` tools.

The `semanage` tool is a policy management tool that modifies the SELinux configuration. It also allows you to configure file contexts, networking component labeling, and user mappings for Linux-to-SELinux.

raw Device Mapping

The `raw` devices interface has been deprecated in Red Hat Enterprise Linux 5; `raw` device mapping is now configured via `udev` rules.

To configure `raw` device mapping, add the appropriate entries to `/etc/udev/rules.d/60-raw.rules` in the following formats:

- For device names:

```
ACTION=="add", KERNEL="<device name>", RUN+="raw /dev/raw/rawX %N"
```

- For major / minor numbers:

```
ACTION=="add", ENV{MAJOR}="A", ENV{MINOR}="B", RUN+="raw /dev/raw/rawX %M %m"
```

Replace `<device name>` with the name of the device you need to bind (for example, `/dev/sda1`). "A" and "B" are the major / minor numbers of the device you need to bind, and `X` is the raw device number that you want the system to use.

If you have a large, pre-existing `/etc/sysconfig/rawdevices` file, convert it with the following script:

```
#!/bin/sh

grep -v "^ *#" /etc/sysconfig/rawdevices | grep -v "^$" | while read dev
major minor ; do
  if [ -z "$minor" ]; then
    echo "ACTION==\"add\", KERNEL==\"${major##*/dev/}\",
RUN+=\"/usr/bin/raw $dev %N\""
  else
    echo "ACTION==\"add\", ENV{MAJOR}==\"$major\", ENV{MINOR}==\"$minor\",
RUN+=\"/usr/bin/raw $dev %M %m\""
  fi
done
```

QLogic Support

Red Hat Enterprise Linux 5 supports the *QLogic* family of iSCSI HBA (*Host Bus Adapters*). At present, only the iSCSI interface to these boards are supported (using the `qla4xxx` driver).

In addition, Red Hat does not currently support these boards as Ethernet NIC, as this capability requires the `qla3xxx` driver. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.

IBM System z Instruction Set

In order to optimally exploit the *IBM System z* instruction set for 31-bit applications, it is recommended that you use the `gcc` option `-march=z900`. For 64-bit applications, the `gcc` will exploit the *IBM System z* instruction set by default.

iSeries Access for Linux

The *iSeries ODBC Driver* for Linux has been replaced by the *iSeries Access for Linux*, which can be downloaded at the following link:

<http://www.ibm.com/eserver/iseries/access/linux/>

The *iSeries Access for Linux* offers Linux-based access to iSeries servers, and allows you to:

- Access the DB2 UDB (*Universal Database*) for iSeries using its ODBC Driver
- Establish a 5250 session to an iSeries server from a Linux client
- Access the DB2 UDB via the EDRS (*Extended Dynamic Remote SQL*) driver
- Support 32-bit (i386 and PowerPC) and 64-bit (x86-64 and PowerPC) platforms

IBM Power4 iSeries

Red Hat Enterprise Linux no longer supports the *IBM Power4 iSeries*.

1.5. Driver Update Program

This section includes information on the implementation of the Red Hat Enterprise Linux 5 Driver Update Program.

Kernel Module Packages

On Red Hat Enterprise Linux 5, it is possible to build updated kernel module packages that depend upon the current kernel ABI version and not on a specific kernel release number. This facilitates building kernel modules that can be used against a range of Red Hat Enterprise Linux 5 kernels, rather than a single release. The project website at <http://www.kerneldrivers.org/> contains more information about the packaging process, as well as several examples.

Note that the following issues have also been identified:

- Bootpath drivers distributed as `kmod` packages are not officially supported.
- Overriding existing in-kernel drivers are not currently supported.

These issues will be addressed in a future update of Red Hat Enterprise Linux 5.

Kernel Module Loading

The module loading behavior on Red Hat Enterprise Linux 5 has changed from previous releases of Red Hat Enterprise Linux. The modules shipped in the Red Hat Enterprise Linux 5 kernel package are signed, as was the case in Red Hat Enterprise Linux 4. On Red Hat Enterprise Linux 5 kernels, however, it is no longer possible to load a signed module from another kernel build.

This means that a module shipped with the initial Red Hat Enterprise Linux 5 distribution cannot be loaded in future updated kernels. This helps prevent users from loading unsupported modules on a system. Red Hat only supports modules that are signed and included in a distribution.

If you want to load an older module, you can try rebuilding it without a signature. Alternatively, you can remove the signature from the binary file using the following command:

```
objcopy -R .module_sig <module name>-mod.ko <module name>-nosig.ko
```


It is recommended that you consult with a designated Red Hat Global Support Representative before attempting to load unsigned modules.

1.6. Internationalization

This section includes information on language support under Red Hat Enterprise Linux 5.

Input Methods

SCIM (Smart Common Input Method) has replaced IIMF as the input method system for Asian and other languages in this release. The default GTK Input Method Module for SCIM is provided by `scim-bridge`; in Qt, it is provided by `scim-qtim`.

Below are the default trigger hotkeys for different languages:

- All languages: **Ctrl-Space**
- Japanese: **Zenkaku-Hankaku** or **Alt-`**
- Korean: **Shift-Space**

If SCIM is installed, it runs by default for all users.

After installing or removing SCIM engine packages, it is recommended to start a new desktop session in order for the changes to be reflected in the SCIM language menu.

Language Installation

To enable additional language support for some Asian languages, you need to install the necessary language support packages. Below is a list of these languages and the command you need to run (as root) to install their corresponding language support packages:

- Assamese – `yum install fonts-bengali m17n-db-assamese scim-m17n`
- Bengali – `yum install fonts-bengali m17n-db-bengali scim-m17n`
- Chinese – `yum install fonts-chinese scim-chewing scim-pinyin scim-tables-chinese`
- Gujarati – `yum install fonts-gujarati m17n-db-gujarati scim-m17n`
- Hindi – `yum install fonts-hindi m17n-db-hindi scim-m17n`
- Japanese – `yum install fonts-japanese scim-anthy`
- Kannada – `yum install fonts-kannada m17n-db-kannada scim-m17n`
- Korean – `yum install fonts-korean scim-hangul`
- Malayalam – `yum install fonts-malayalam m17n-db-malayalam scim-m17n`
- Marathi – `yum install fonts-hindi m17n-db-marathi scim-m17n`
- Oriya – `yum install fonts-oriya m17n-db-oriya scim-m17n`
- Punjabi – `yum install fonts-punjabi m17n-db-punjabi scim-m17n`

- Sinhala – `yum install fonts-sinhala m17n-db-sinhala scim-m17n`
- Tamil – `yum install fonts-tamil m17n-db-tamil scim-m17n`
- Telugu – `yum install fonts-telugu m17n-db-telugu scim-m17n`

It is also recommended that you install `scim-bridge-gtk` and `scim-qtimm` when enabling additional language support. The `scim-bridge-gtk` package prevents possible binary conflicts with third-party applications linked against older versions of `libstdc++`.

Note that additional language support packs are also available for OpenOffice (`openoffice.org-langpack-<language code>_<locale>`) and KDE (`kde-i18n-<language>`). These packages can also be installed through `yum`.

im-chooser

A new user configuration tool called `im-chooser` has been added, which allows you to easily disable or enable the usage of input methods on your desktop. So if SCIM is installed but you do not wish to run it on your desktop, you can disable it using `im-chooser`.

xinputrc

At X startup, `xinput.sh` now sources `~/.xinputrc` or `/etc/X11/xinit/xinputrc` instead of searching config files under `~/.xinput.d/` or `/etc/xinit/xinput.d/`.

Pango Support in Firefox

Firefox in Red Hat Enterprise Linux 5 is built with Pango, which provides better support for certain scripts, such as Indic and some CJK scripts.

To disable the use of Pango, set `MOZ_DISABLE_PANGO=1` in your environment before launching Firefox.

Fonts

Support is now available for synthetic emboldening of fonts that do not have a bold face.

New fonts for Chinese have been added: AR PL ShanHeiSun Uni (`uming.ttf`) and AR PL ZenKai Uni (`ukai.ttf`). The default font is AR PL ShanHeiSun Uni, which contains embedded bitmaps. If you prefer outline glyphs, add the following section to your `~/.font.conf` file:

```
<fontconfig>
  <match target="font">
    <test name="family" compare="eq">
      <string>AR PL ShanHeiSun Uni</string>
    </test>
    <edit name="embeddedbitmap" mode="assign">
      <bool>>false</bool>
    </edit>
  </match>
</fontconfig>
```

gtk2 IM submenu

The Gtk2 context menu IM submenu no longer appears by default. You can enable it on the command line with the following command:

```
gconftool-2 --type bool --set  
'/desktop/gnome/interface/show_input_method_menu' true
```

Support for text installation on CJK

CJK (Chinese, Japanese, and Korean) rendering support has been removed from the Anaconda text installation. The text installation method is being deprecated in the long term, as the GUI installation, VNC and kickstart methods are preferred.

gtk+ deprecation

The following packages are deprecated and scheduled for removal in Red Hat Enterprise Linux:

- **gtk+**
- **gdk-pixbuf**
- **glib**

These packages are being deprecated in favor of the **gtk2** stack, which offers better functionality particularly in terms of internationalization and font handling.

CJK input on console

If you need to display Chinese, Japanese, or Korean text on the console, you need to set up a framebuffer; afterwards, install **bogl-bterm**, and run **bterm** on the framebuffer.

1.7. Kernel Notes

This section notes the differences between 2.6.9 (on which Red Hat Enterprise Linux 4 is based) and 2.6.18 (which Red Hat Enterprise Linux 5 will inherit) as of July 12, 2006. Additional features which we are currently working on upstream (for example, virtualization) that will appear late in 2.6.18 or 2.6.19 are not highlighted here. In other words, this list only shows what is already included in the upstream Linus tree; not what is currently in development. Consequently, this list is not a final, or complete list of the new Red Hat Enterprise Linux 5 features, although it does give a good overview of what can be expected. Also, note that this section only picks out highlights of upstream changes, and as such it is not fully comprehensive. It does not include mention of several low-level hardware support enhancements and device driver info.

The following is a good source for a next level-of-detail view:

<http://kernelnewbies.org/LinuxChanges>

Performance / Scalability

- Big Kernel Lock preemption (2.6.10)
- Voluntary preemption patches (2.6.13) (subset in Red Hat Enterprise Linux 4)
- Lightweight user-space priority inheritance (PI) support for futexes, useful for real-time applications (2.6.18)
 - ref: <http://lwn.net/Articles/178253/>
- New 'mutex' locking primitive (2.6.16)

- High resolution timers (2.6.16)
 - In contrast to the low-resolution timeout API implemented in kernel/timer.c, hrtimers provide finer resolution and accuracy depending on system configuration and capabilities. These timers are currently used for itimers, POSIX timers, nanosleep and precise in-kernel timing.
- Modular, on-the-fly switchable I/O schedulers (2.6.10)
 - This was adjustable only by boot option in Red Hat Enterprise Linux 4 (also system-wide instead of per-queue).
- Conversion to 4-level page tables (2.6.11)
 - allows x86-64 to increase from 512G to 128TB of memory
- New Pipe implementation (2.6.11)
 - 30-90% performance improvement in pipe bandwidth
 - circular buffer allows more buffering than blocking writers
- "Big Kernel Semaphore": turns the Big Kernel Lock into a semaphore
 - reduces latency by breaking up long lock hold times and adding voluntary preemption
- X86 "SMP alternatives"
 - optimizes a single kernel image at runtime according to the available platform
 - ref: <http://lwn.net/Articles/164121/>
- libhugetlbfs
 - allows applications to use the huge page support in Linux with no need for source code modification
- kernel-headers package
 - replaces the `glibc-kernheaders` package
 - provides better suitability with the new `headers_install` feature of the 2.6.18 kernel
 - notable kernel header-related changes:
 - removed `<linux/compiler.h>` header file, as it is no longer useful
 - removed `__syscallX()` macros; user-space should use `syscall()` from the C library instead
 - removed `<asm/atomic.h>` and `<asm/bitops.h>` header files; C compiler provides its own atomic built-in functions better suitable for user-space programs
 - content previously protected with `#ifdef __KERNEL__` is now removed completely with the `unifdef` tool; defining `__KERNEL__` in order to view parts which should not be visible to user-space is no longer effective

- removed the `PAGE_SIZE` macro from some architectures, due to variance in page sizes; user-space should be using `sysconf (_SC_PAGE_SIZE)` or `getpagesize()`
- to provide better suitability for user-space, removed several header files and header content

Generic Feature Additions

- `kexec` and `kdump` (2.6.13)
 - `diskdump` and `netdump` have been replaced by `kexec` and `kdump`, which ensure faster boot-up and creation of reliable kernel vmcores for diagnostic purposes. For more information and configuration instructions, please refer to `/usr/share/doc/kexec-tools-<version>/kexec-kdump-howto.txt` (replace `<version>` with the corresponding version of the `kexec-tools` package installed).
 - Note that at present, virtualized kernels cannot use the `kdump` function.
- `inotify` (2.6.13)
 - user interface for this is through the following syscalls: `sys_inotify_init`, `sys_inotify_add_watch`, and `sys_inotify_rm_watch`.
- Process Events Connector (2.6.15)
 - reports `fork`, `exec`, id change, and exit events for all processes to user-space.
 - Applications that may find these events useful include accounting / auditing (for example, ELSA), system activity monitoring (for example, `top`), security, and resource management (for example, CKRM). Semantics provide the building blocks for features like per-user-namespace, "files as directories" and versioned file systems.
- Generic RTC (RealTime Clock) subsystem (2.6.17)
- `splice` (2.6.17)
 - new IO mechanism which avoids data copies when transferring data between applications
 - ref: <http://lwn.net/Articles/178199/>

File System / LVM

- EXT3
 - support for Extended Attributes in the body of large inode in ext3: saves space and improves performance in some cases (2.6.11)
- Device mapper multipath support
- ACL support for NFSv3 and NFSv4 (2.6.13)
- NFS: supports large reads and writes on the wire (2.6.16)
 - The Linux NFS client now supports transfer sizes of up to 1MB.

- VFS changes
 - The "shared subtree" patches have been merged. (2.6.15)
 - ref: <http://lwn.net/Articles/159077/>
- Big CIFS update (2.6.15)
 - features several performance improvements as well as support for Kerberos and CIFS ACL
- autofs4: updated to provide direct mount support for user-space autofs (2.6.18)
- cachefs core enablers (2.6.18)

Security


- Multilevel security implementation for SELinux (2.6.12)
- Audit subsystem
 - support for process-context based filtering (2.6.17)
 - more filter rule comparators (2.6.17)
- TCP/UDP getpeercon: enabled security-aware applications to retrieve the entire security context of a process on the other side of a socket using an IPSec security association. If only MLS-level information is needed or interoperability with legacy unix system is required, NetLabel can be used in place of IPSec.

Networking

- Added several TCP congestion modules (2.6.13)
- IPv6: supports several new sockopt / ancillary data in Advanced API (2.6.14)
- IPv4/IPv6: UFO (UDP Fragmentation Offload) Scatter-gather approach (2.6.15)
 - UFO is a feature wherein the Linux kernel network stack will offload the IP fragmentation functionality of large UDP datagram to hardware. This will reduce the overhead of stack in fragmenting the large UDP datagram to MTU-sized packets.
- Added nf_contrack subsystem (2.6.15)
 - The existing connection tracking subsystem in netfilter can only handle ipv4. There were two choices present to add connection tracking support for ipv6; either duplicate all of the ipv4 connection tracking code into an ipv6 counterpart, or (the choice taken by these patches) design a generic layer that could handle both ipv4 and ipv6 and thus requiring only one sub-protocol (TCP, UDP, etc.) connection tracking helper module to be written. In fact, nf_contrack is capable of working with any layer 3 protocol.
- IPV6
 - RFC 3484-compliant source address selection (2.6.15)
 - added support for Router Preference (RFC4191) (2.6.17)
 - added Router Reachability Probing (RFC4191) (2.6.17)

- added support for Multiple Routing Tables and Policy Routing
- Wireless updates
 - hardware crypto and fragmentation offload support
 - QoS (WME) support, "wireless spy support"
 - mixed PTK/GTK
 - CCMP/TKIP support and WE-19 HostAP support
 - BCM43xx wireless driver
 - ZD1211 wireless driver
 - WE-20, version 20 of the Wireless Extensions (2.6.17)
 - added the hardware-independent software MAC layer, "Soft MAC" (2.6.17)
 - added LEAP authentication type
- Added generic segmentation offload (GSO) (2.6.18)
 - can improve performance in some cases, though it needs to be enabled through `ethtool`
- DCCPv6 (2.6.16)

Added Hardware Support



NOTE

This section only enumerates the most generic features among many.

- x86-64 clustered APIC support (2.6.10)
- Infiniband support (2.6.11)
- Hot plug
 - added generic memory add/remove and supporting functions for memory hotplug (2.6.15)
 - hot plug CPU support for physically adding new processors (hotplug disable / enable of already existing CPUs is already supported)
- SATA/libata enhancements, additional hardware support
 - A completely reworked libata error handler; the result of all this work should be a more robust SATA subsystem which can recover from a wider range of errors.
 - Native Command Queuing (NCQ), the SATA version of tagged command queuing - the ability to have several I/O requests to the same drive outstanding at the same time. (2.6.18)
 - Hotplug support (2.6.18)

- EDAC support (2.6.16)
 - The EDAC goal is to detect and report errors that occur within the system.
- Added a new ioatdma driver for the Intel(R) I/OAT DMA engine (2.6.18)


NUMA (Non-Uniform Memory Access) / Multi-core

- Cpusets (2.6.12)
 - Cpusets now provide a mechanism for assigning a set of CPUs and Memory Nodes to a set of tasks. Cpusets constrain the CPU and memory placement of tasks only to the resources within a task's current cpuset. These are essential in managing dynamic job placement on large systems.
- NUMA-aware slab allocator (2.6.14)
 - This creates slabs on multiple nodes and manages slabs in such a way that locality of allocations is optimized. Each node has its own list of partial, free and full slabs. All object allocations for a node occur from node-specific slab lists.
- Swap migration (2.6.16)
 - Swap migration allows the moving of physical location of pages between nodes in a NUMA system while the process is running.
- Huge pages (2.6.16)
 - Added NUMA policy support for huge pages: the `huge_zone_list()` function in the memory policy layer provides a list of zones ordered by NUMA distance. The `hugetlb` layer will walk that list looking for a zone that has available huge pages but is also in the nodeset of the current cpuset.
 - Huge pages now obey cpusets.
- Per-zone VM counters
 - provide zone-based VM statistics, which are necessary in determining what state of memory a zone is in
- Netfilter ip_tables: NUMA-aware allocation. (2.6.16)
- Multi-core
 - Added a new scheduler domain for representing multi-core with shared caches between cores. This makes it possible to make smarter cpu scheduling decisions on such systems, improving performance greatly for some cases. (2.6.17)
 - Power saving policy for the CPU scheduler: with multicore/smt cpus, the power consumption can be improved by leaving some packages idle while others do all the work, instead of spreading the tasks over all CPUs.

2. RELEASE NOTES FOR X86_64

2.1. Installation-Related Notes

The following section includes information specific to the installation of Red Hat Enterprise Linux and the Anaconda installation program.



NOTE

In order to upgrade an already-installed Red Hat Enterprise Linux, you must use Red Hat Network to update those packages that have changed.

You may use Anaconda to perform a fresh installation of Red Hat Enterprise Linux 5 or to perform an upgrade from *the latest updated version* of Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5.

If you are copying the contents of the Red Hat Enterprise Linux 5 CD-ROMs (in preparation for a network-based installation, for example) be sure to copy the CD-ROMs for the operating system *only*. Do not copy the Supplementary CD-ROM, or any of the layered product CD-ROMs, as this will overwrite files necessary for Anaconda's proper operation. These CD-ROMs must be installed *after* Red Hat Enterprise Linux has been installed.

2.1.1. ISO Contents and Registration

With optional content in the same tree or ISO image, it is important to avoid a mismatch between the components offered for installation and those covered by the subscription. Such a mismatch could result in an increased exposure to bug and vulnerability risks.

In order to ensure that the components offered are in sync with the subscription, Red Hat Enterprise Linux 5 requires entering an *Installation Number* that will be used to configure the installer to offer the right package set. This Installation Number is included in your subscription.

If you skip entering the Installation Number, this will result in a core Server or Desktop installation. Additional functionality can then be added manually at a later time. For more information about Installation Numbers, please refer to <http://www.redhat.com/apps/support/in.html>.

The Installation Number used during the installation process will be saved in `/etc/sysconfig/rhn/install-num`. When registering with Red Hat Network, this file will be referenced by `rhn_register` to automatically determine which appropriate child channels the system should be subscribed to.

2.1.2. New RPM GPG Signing Keys

A new release signing key is used to sign Red Hat Enterprise Linux 5 packages. When updating a system for the first time, you will be prompted to allow this key to be installed.

Signing keys are distributed in the following files:

- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release` – contains the public key for the new release signing key
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-auxiliary` – contains the public key for an auxiliary release signing key, currently not in use
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-former` – contains the public key for the previous release signing key, used for past Red Hat Enterprise Linux releases

2.1.3. Subversion

In Red Hat Enterprise Linux 5, the Subversion version control system is linked against Berkeley DB 4.3. If you are upgrading from Red Hat Enterprise Linux 4 and any Subversion repositories which use the Berkeley DB backend "BDB" (rather than the pure file system-based "FSFS" backend) have been created on the system, special care must be taken to ensure the repositories can be accessible after the upgrade. The following process must be performed on the Red Hat Enterprise Linux 4 system, *prior* to upgrading to Red Hat Enterprise Linux 5:

1. Shut down any running processes and ensure that no processes can access the repository (for example, `httpd`, `svnserve` or any local users with direct access).
2. Create a backup of the repository using the following command:

```
svnadmin dump /path/to/repository | gzip > repository-backup.gz
```

3. Run the `svnadmin recover` command on the repository:

```
svnadmin recover /path/to/repository
```

4. Delete any unused log files in the repository:

```
svnadmin list-unused-dblogs /path/to/repository | xargs rm -vf
```

5. Delete any remaining shared-memory files in the repository:

```
rm -f /path/to/repository/db/__db.0*
```

2.1.4. Other Installation Notes

- If IDE/PATA (*Parallel ATA*) devices are configured in "100% Native" mode, some BIOSes may prevent the Red Hat Enterprise Linux 5 installation process from completing successfully. To prevent this from occurring, configure the IDE/PATA mode as "Legacy" in the BIOS.
- The *IBM System z* does not provide a traditional Unix-style physical console. As such, Red Hat Enterprise Linux 5 for the *IBM System z* does not support the `firstboot` functionality during initial program load.

To properly initialize setup for Red Hat Enterprise Linux 5 on the *IBM System z*, run the following commands after installation:

- `/usr/bin/setup` – provided by the `setuptools` package
- `/usr/bin/rhn_register` – provided by the `rhn-setup` package
- When booting Anaconda with PXE using the parameter `ksdevice=bootif`, you will still be prompted for the ethernet interface to use during installation. If only one ethernet device is plugged in, use the parameter `ksdevice=link` instead. Alternatively, you can also specify the interface manually.

2.2. Technology Previews

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 5 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. Erratas 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.

Stateless Linux

Included in this release of Red Hat Enterprise Linux 5 are enabling infrastructure pieces for *Stateless Linux*. Stateless Linux is a new way of thinking about how a system is to be run and managed, designed to simplify provisioning and management of large numbers of systems by making them easily replaceable. This is accomplished primarily by establishing prepared system images which get replicated and managed across a large number of stateless systems, running the operating system in a read-only manner (please refer to `/etc/sysconfig/readonly-root` for more details).

In its current state of development, the Stateless features are subsets of the intended goals. As such, the capability is being labeled as a technology preview.

The following is a list of the initial capabilities included in Red Hat Enterprise Linux 5:

- running a stateless image over NFS
- running a stateless image via loopback over NFS
- running on iSCSI

It is highly recommended that those interested in testing stateless code read the HOWTO at <http://fedoraproject.org/wiki/StatelessLinuxHOWTO> and join stateless-list@redhat.com.

GFS2

GFS2 is an evolutionary advancement based on the GFS file system. While fully functional, *GFS2* is not yet considered production-ready. *GFS2* is targeted to become fully supported in a subsequent Red Hat Enterprise Linux 5 update. There is also an in-place conversion utility, `gfs2_convert`, which can update the metadata of the older GFS file system format, converting it to a *GFS2* file system.

FS-Cache

FS-Cache is a local caching facility for remote file systems that allows users to cache NFS data on a locally mounted disk. To set up the *FS-Cache* facility, install the `cachefilesd` RPM and refer to the instructions in `/usr/share/doc/cachefilesd-<version>/README`.

Replace `<version>` with the corresponding version of the `cachefilesd` package installed.

Compiz

Compiz is an OpenGL-based compositing window manager. In addition to regular window management, `compiz` also acts as a compositing manager, coordinating and synchronizing the overall desktop redrawing to provide a smoother desktop experience with less flicker.

`Compiz` uses 3D hardware acceleration to render effects such as live thumbnail windows, window drop shadows, animated window minimizing and transitions between virtual desktops.

Due to limitations in the current rendering architecture, `compiz` cannot work correctly with direct rendering OpenGL applications or applications using the Xv extension. Such applications will exhibit harmless rendering artifacts; because of this, `compiz` is currently a technology preview.

Enhancement for Ext3

In Red Hat Enterprise Linux 5, the EXT3 file system capacity has been extended beyond 8TB to a maximum of 16TB. This capability is being included as a technology preview, and is targeted for full support in a future release of Red Hat Enterprise Linux 5.

AIGLX

AIGLX is a technology preview feature of the otherwise fully supported X server. It aims to enable GL-accelerated effects on a standard desktop. The project consists of the following:

- a lightly modified X server
- an updated Mesa package that adds new protocol support

By installing these components, you can have GL-accelerated effects on your desktop with very few changes, as well as the ability to enable and disable them at will without replacing your X server. *AIGLX* also enables remote GLX applications to take advantage of hardware GLX acceleration.

Frysk GUI

The goal of the *frysk* project is to create an intelligent, distributed, always-on system monitoring and debugging tool that allows developers and system administrators to:

- monitor running processes and threads (including creation and destruction events)
- monitor the use of locking primitives
- expose deadlocks
- gather data
- debug any given process by choosing it from a list or allowing *frysk* to open a source code (or other) window on a process that is crashing or misbehaving

In Red Hat Enterprise Linux 5 the *frysk* graphical user interface is a technology preview, whereas the *frysk* command line interface is fully supported.

Systemtap

Systemtap provides free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. With the help of *systemtap*, developers no longer need to go through the tedious and disruptive instrument, recompile, install, and reboot sequence that may be otherwise required to collect data.

Dogtail

Dogtail is a GUI test tool and automation framework written in Python that uses *Accessibility* technologies to communicate with desktop applications.

Support for Indic Languages and Sinhalese

Red Hat Enterprise Linux 5 also features support for the following languages as technology preview:

- Assamese

- *Assamese*
- *Kannada*
- *Sinhalese*
- *Telugu*

For more information about how to install and enable support for these languages, refer to the *Internationalization* section of this document.

Installing to dm-multipath Devices

Anaconda now has the capability to detect, create, and install to *dm-multipath* devices. To enable this feature, add the parameter `mpath` to the kernel boot line.

Note that the parameter `mpath` may cause a boot failure if a device's `major:minor` number changes. This issue will be addressed in a future update of Red Hat Enterprise Linux 5.

Installation / Boot for iSCSI software initiator (open-iscsi)

Anaconda now provides the ability to install to an iSCSI device. Booting and installing is fully supported with the *QLogic qla4xxx* hardware initiator. However, the capability to install to an iSCSI device for the `open-iscsi` software initiator is currently considered a Technology Preview, due to the following issues:

- Text mode installation does not complete. You must do a graphical install, or an automated kickstart install.
- Media-based installations do not complete. You must do a network-based install.
- Depending on the timing of events, Anaconda may be unable to detect all the iSCSI targets or LUNs. When this occurs, use the installer shell to configure the storage through iSCSI commands.
- The `iscsid` daemon may not properly start. Such an occurrence will prevent the system from handling all iSCSI errors, such as network problems, SCSI/iSCSI timeouts, and target errors. To confirm that the `iscsid` daemon is running, run the command `iscsiadm -m session -i` and check that the line, **Internal iscsid Session State:** reports a value (it can be any value).
- On certain iSCSI target implementations, the system may hang during shutdown.
- On certain iSCSI target implementations, the system may hang during reboot. To avoid this, shutdown the system and boot it up again (instead of rebooting directly from a session).
- Booting from iSCSI devices on the *IBM System p* does not work reliably. While installation on an iSCSI device may appear to succeed, the resulting installation will not boot properly.
- On the first boot after install, you may receive SELinux errors such as the following:

```
kernel: audit(1169664832.270:4): avc: denied { read
} for pid=1964 comm="iscsid"
```

To work around this, boot the system with the kernel parameter `enforcing=0`. Once the system has properly booted, use the command `setenforce 1` to restore enforcing mode.

These limitations will be addressed in a future Red Hat Enterprise Linux 5 update.

2.3. Known Issues

- Host bus adapters that use the MegaRAID driver must be set to operate in "Mass Storage" emulation mode, not in "I2O" emulation mode. To do this, perform the following steps:

1. Enter the **MegaRAID BIOS Set Up Utility**.
2. Enter the **Adapter settings** menu.
3. Under **Other Adapter Options**, select **Emulation** and set it to **Mass Storage**.

If the adapter is incorrectly set to "I2O" emulation, the system will attempt to load the i2o driver. This will fail, and prevent the proper driver from being loaded.

Previous Red Hat Enterprise Linux releases generally do not attempt to load the I2O driver before the MegaRAID driver. Regardless of this, the hardware should always be set to "Mass Storage" emulation mode when used with Linux.

- Red Hat Enterprise Linux 5 includes **openmpi-1.1.1-4.el5** (from the OFED 1.1 distribution), which has been discovered to eventually quit working entirely. This happens after the **openmpi** stack works as expected for a varying amount of time.

For updated versions of **openmpi**, please check <http://people.redhat.com/dledford/Infiniband/openmpi>

- Installing Windows Server 2003 as a guest on a fully virtualized Red Hat Enterprise Linux 5 system ends unexpectedly after completing the first stage of installation. When this happens, the graphical console window closes, and the guest disappears from the **Virtual Machine Manager's** list of machines, resulting in a **Broken pipe** error.

This issue will be resolved in an upcoming Red Hat Enterprise Linux 5 update. To work around this, use the following command at the terminal:

```
xm create /etc/xen/<name of guest machine>
```

Afterwards, open the virtual machine.

- When attempting to create a fully virtualized Windows Server 2003 from a CD / DVD, the second stage of the guest install will not continue upon reboot.

To work around this, edit **/etc/xen/<name of guest machine>** by properly appending an entry for the CD / DVD device.

If an installation to a simple file is used as a virtual device, the **disk** line of **/etc/xen/<name of guest machine>** will read like the following:

```
disk = [ 'file:/PATH-OF-SIMPLE-FILE,hda,w' ]
```

A DVD-ROM device located on the host as **/dev/dvd** can be made available to stage 2 of the installation as **hdc** by appending an entry like **'phy:/dev/dvd,hdc:cdrom,r'**. As such, the disk line should now read as follows:

```
disk = [ 'file:/opt/win2003-sp1-20061107,hda,w',  
        'phy:/dev/dvd,hdc:cdrom,r' ]
```

The precise device path to use may vary depending on your hardware.

- `rmmmod xennet` causes `domU` to crash; this is caused by a grant table issue in the Virtualization feature. Due to the current inability of the Virtualization feature to asynchronously release grant table operations, it is unsafe to unload the `xennet` module in guests. In such situations, grant tables are used to perform backend-frontend communication, and there is no guarantee that the backend will release the references, leading to an inevitable memory leak.

This issue will be resolved in the next minor release of Red Hat Enterprise Linux 5. At present, users are advised *not* to unload the `xennet` module in guests.

- Running `ethtool eth0` outputs incomplete information about the ethernet card settings. This only occurs in systems running a virtualized kernel, since the Virtualization feature uses a networking setup where the physical ethernet device is identified as `peth0`. As such, the correct command for retrieving information about the physical ethernet device is `ethtool peth0`.
- When using Red Hat Enterprise Linux 5 on a machine with an *nVidia CK804* chipset installed, you may receive kernel messages similar to the following:

```
kernel: assign_interrupt_mode Found MSI capability
kernel: pcie_portdrv_probe->Dev[005d:10de] has invalid IRQ. Check
vendor BIOS
```

These messages indicate that certain PCI-E ports are not requesting IRQs. Further, these messages do not, in any way, affect the operation of the machine.

- Some *Cisco Aironet Wireless* devices prevent **NetworkManager** from storing connection details for wireless networks that do not broadcast an SSID. This is caused by a *Cisco Aironet Wireless* device firmware limitation.
- Laptops that have the *Cisco Aironet MPI-350* wireless card equipped may hang trying to get a DHCP address during any network-based installation using the wired ethernet port.

To work around this, use local media for your installation. Alternatively, you can disable the wireless card in the laptop BIOS prior to installation (you can re-enable the wireless card after completing the installation).

- Currently, `system-config-kickstart` does not support package selection and deselection. When using `system-config-kickstart`, the **Package Selection** option indicates that it is disabled. This is because `system-config-kickstart` uses `yum` to gather group information, but is unable to configure `yum` to connect to Red Hat Network.

This issue is currently being investigated for resolution by the next minor release of Red Hat Enterprise Linux 5. At present, you need to update package sections in your kickstart files manually. When using `system-config-kickstart` to open a kickstart file, it will preserve all package information in it and write it back out when you save.

- 4-socket AMD *Sun Blade X8400 Server Modules* systems that do not have memory configured in `node 0` will panic during boot. Systems should be configured with memory in `node 0` to prevent the kernel panic.
- Installing to LVM mirror devices through Anaconda is currently not supported. This capability will be added in a future update of Red Hat Enterprise Linux 5.

- When installing Red Hat Enterprise Linux 5 from a directory on an NFS server containing Red Hat Enterprise Linux ISO images, Anaconda may display the following error message:

```
Unable to read package metadata. This may be due to a missing
repopdata directory.
Please ensure that your install tree has been correctly generated.
Cannot
open/read repomd.xml file for repository:
```

This problem occurs if the directory holding the ISO images also contains a partially unpacked installation tree (for example, the `/images` directory from the first ISO). The presence of such directories results in the error stated above.

To prevent this error, unpack trees only to directories other than the one containing the installation ISO images.

- Boot-time logging to `/var/log/boot.log` is not available in this release of Red Hat Enterprise Linux 5. An equivalent functionality will be added in a future update of Red Hat Enterprise Linux 5.

- Neither `kexec` nor `kdump` are able to dump onto disks attached to an *accraid* controller.

To work around this issue, use `scp` for network dumping. Alternatively, you can also dump onto a disk through a different controller.

- Installing a fully virtualized guest using split installation media -- specifically, multiple CD-ROMs -- may fail when required to switch between installation CDs. During the guest OS installation process, users may be prevented from mounting or ejecting installation CDs, which prevents the installation from completing.

As such, it is recommended that you use the QEMU monitor console to switch CD-ROM images during the guest OS installation process. The procedure is as follows:

1. Open a graphical VNC console to the guest OS.
2. Unmount the CD-ROM device in the guest OS.
3. Switch to the QEMU monitor console by pressing `Ctrl-Alt-2`.
4. Run the command `eject hdc`.
5. Run the command `change hdc <path to the CD-ROM in host system>`.
6. Switch back to the guest OS console by pressing `Ctrl-Alt-1`.
7. Mount the CD-ROM device in the guest OS.

Note that when using a regular VNC client the host X server may encounter some difficulty interpreting the `Ctrl-Alt-2` and `Ctrl-Alt-1` command. To work around this in `virt-manager`, use *sticky keys*. Pressing `Ctrl` three times makes it "sticky" until the next non-modifier is pressed. As such, to send `Ctrl-Alt-1`, press `Ctrl` twice before pressing `Ctrl-Alt-1`.

- The Red Hat Enterprise Linux 5 *Driver Update Model* creates modified `initrd` images whenever a `kmod` package that includes a bootpath-modifying driver is installed. In time, the number of backup `initrd` images may soon fill the `/boot` partition, particularly if the system

undergoes a sizeable number of driver updates.

As such, it is recommended that you monitor the free space on the `/boot` partition if you regularly perform driver updates. You can free up more space in `/boot` by removing older `initrd` images; these files end in `.img0`, `.img1`, `.img2`, and so on.

- Autorun on removable media is currently disabled. To install packages from the Red Hat Enterprise Linux Supplementary CD, launch the CD installer manually using the following command:

```
system-cdinstall-helper /media/path-to-mounted-drive
```

- When upgrading from Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5, the Deployment Guide is not automatically installed. You need to use `pirut` to manually install it after completing the upgrade.
- An `autofs` bug prevents multi-mounts from working properly.

During an expiry, if the last multi-mount component to be checked does not have a mount associated with it while other components are busy, `autofs` erroneously determines the multi-mount to be expirable. This causes the multi-mount to be partially expired, resulting in the multi-mount becoming unresponsive to further mount requests and expire runs.

To permanently resolve this problem, update `autofs` using the command `yum update autofs`.

- The system may not successfully reboot into a `kexec/kdump` kernel if X is running and using a driver other than `vesa`. This problem only exists with *ATI Rage XL* graphics chipsets.

If X is running on a system equipped with *ATI Rage XL*, ensure that it is using the `vesa` driver in order to successfully reboot into a `kexec/kdump` kernel.

2.4. General Information

This section contains general information not specific to any other section of this document.

Red Hat Enterprise Linux Deployment Guide

This release of Red Hat Enterprise Linux includes a fairly comprehensive Deployment Guide. To access it, go to **System** (on the top panel) => **Documentation** => **Red Hat Enterprise Linux Deployment Guide**.

It is the intention of Red Hat to provide fully localized versions of the Deployment Guide for all supported languages. If you have installed a localized version of the Deployment Guide, it is recommended that you update it when a new version becomes available through Red Hat Network.

Web Server Packaging Changes

Red Hat Enterprise Linux 5 now includes version 2.2 of the Apache HTTP Server. This release brings a number of improvements over the 2.0 series, including:

- improved caching modules (`mod_cache`, `mod_disk_cache`, `mod_mem_cache`)
- a new structure for authentication and authorization support, replacing the authentication modules provided in previous versions
- support for proxy load balancing (`mod_proxy_balancer`)

- support for handling large files (namely, greater than 2GB) on 32-bit platforms

The following changes have been made to the default httpd configuration:

- The `mod_cern_meta` and `mod_asis` modules are no longer loaded by default.
- The `mod_ext_filter` module is now loaded by default.

If you are upgrading from a previous release of Red Hat Enterprise Linux, the httpd configuration will need to be updated for httpd 2.2. For more information, refer to <http://httpd.apache.org/docs/2.2/upgrading.html>.

Note that any third-party modules compiled for httpd 2.0 must be rebuilt for httpd 2.2.

php

Version 5.1 of PHP is now included in Red Hat Enterprise Linux 5, which includes a number of changes to the language along with significant performance improvements. Some scripts might need to be edited for use with the new version; please refer to the link below for more information on migrating from PHP 4.3 to PHP 5.1:

<http://www.php.net/manual/en/migration5.php>

The `/usr/bin/php` executable is now built using the CLI command-line SAPI, rather than the CGI SAPI. Use `/usr/bin/php-cgi` for CGI SAPI. The `php-cgi` executable also includes FastCGI support.

The following extension modules have been added:

- the `mysqli` extension, a new interface designed specifically for MySQL 4.1 (included in the `php-mysql` package)
- `date`, `hash`, `Reflection`, `SPL` and `SimpleXML` (built-in with the `php` package)
- `pdo` and `pdo_sqlite` (in the `php-pdo` package)
- `pdo_mysql` (in the `php-mysql` package)
- `pdo_pgsql` (in the `php-pgsql` package)
- `pdo_odbc` (in the `php-odbc` package)
- `soap` (in the `php-soap` package)
- `xmlreader` and `xmlwriter` (in the `php-xml` package)
- `dom` (replacing the `domxml` extension in the `php-xml` package)

The following extension modules are no longer included:

- `dbx`
- `dio`
- `yp`
- `overload`

- `domxml`

The PEAR Framework

The PEAR framework is now packaged in the `php-pear` package. Only the following PEAR components are included in Red Hat Enterprise Linux 5:

- `Archive_Tar`
- `Console_Getopt`
- `XML_RPC`

Encrypted Swap Partitions and Non-root File Systems

Red Hat Enterprise Linux 5 now provides basic support for encrypted swap partitions and non-root file systems. To use these features, add the appropriate entries to `/etc/crypttab` and reference the created devices in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry:

```
my_swap /dev/hdb1 /dev/urandom swap,cipher=aes-cbc-essiv:sha256
```

This creates the encrypted block device `/dev/mapper/my_swap`, which can be referenced in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry for a file system volume:

```
my_volume /dev/hda5 /etc/volume_key cipher=aes-cbc-essiv:sha256
```

The `/etc/volume_key` file contains a plaintext encryption key. You can also specify `none` as the key file name; this configures the system to ask for the encryption key during boot instead.

It is recommended to use LUKS (Linux Unified Key Setup) for setting up file system volumes. To do this, follow these steps:

1. Create the encrypted volume using `cryptsetup luksFormat`.
2. Add the necessary entry to `/etc/crypttab`.
3. Set up the volume manually using `cryptsetup luksOpen` (or reboot).
4. Create a file system on the encrypted volume.
5. Add the necessary entry to `/etc/fstab`.

mount and umount

The `mount` and `umount` commands no longer directly support NFS; a built-in NFS client no longer exists. A separate `nfs-utils` package, which provides `/sbin/mount.nfs` and `/sbin/umount.nfs` helpers, must be installed for this.

CUPS Printer Browsing

CUPS printer browsing over a local subnet can be configured using the graphical tool `system-config-printer`. It can also be done using the CUPS web interface, <http://localhost:631/>.

To use directed broadcasts for printer browsing between subnets, open `/etc/cups/cupsd.conf` on the clients and replace `BrowseAllow @LOCAL` with `BrowseAllow ALL`.

ATI and R500 Support

ATI graphics cards based on the R500 chipset are supported for the `vesa` driver only, and are not supported by Red Hat Enterprise Linux 5 on external monitors, LCD projectors or accelerated 3D support.

up2date and yum

`up2date` is being deprecated in favor of `yum` (Yellowdog Updater Modified). As such, it is advisable that you revise any `up2date`-dependent scripts your system is using accordingly. For more information about `yum`, consult its man page with the command `man yum`; you can also consult the installed documentation under the directories `/usr/share/doc/yum-<version>` and `/usr/share/doc/yum-metadata-parser-<version>` (replace `<version>` with the corresponding version of `yum` and `yum-metadata-parser` installed).

OpenLDAP Server and Red Hat Directory Server

Red Hat Directory Server is an LDAP-based server that centralizes enterprise and network data into an OS-independent, network-based registry. It is set to replace OpenLDAP server components, which will be deprecated *after* Red Hat Enterprise Linux 5. For more information about Red Hat Directory Server, refer to <http://www.redhat.com/software/rha/directory/>.

i810 Driver and i830 Support

The `i810` driver supports all integrated Intel graphics chipsets, from `i810` to `i965`. However, the support for `i830` (and newer) chipsets is limited; the `i810` driver can only set modes listed in the video BIOS. If your machine has an `i830` or newer chipset installed, run the following command to determine what the available modes are:

```
grep Mode: /var/log/Xorg.0.log
```

Modes marked with an asterisk (*) are available for selection.

Many laptop video BIOSes do not supply a mode that matches the native panel size. Therefore the chosen mode may appear stretched, distorted, or with black borders. As such, if your chosen mode does not display properly, you need a BIOS update from your hardware vendor for the native panel size to work correctly.

Intel PRO/Wireless 3945ABG Network Connection Support

This release of Red Hat Enterprise Linux 5 includes support for the `ipw3945` (*Intel PRO/Wireless 3945ABG Network Connection*) adapter. The Red Hat Enterprise Linux 5 Supplementary disc contains the driver, regulatory daemon and firmware needed to support this adapter.

To enable support for the `ipw3945` wireless adapter, search the Red Hat Enterprise Linux 5 Supplementary disc for packages with filenames containing "3945" and install them.

rawio

`rawio` is a deprecated interface; however, Red Hat Enterprise Linux 5 still includes support for it. If you have an application that performs device access using `rawio`, it is highly recommended that you modify your application to open the block device with the `O_DIRECT` flag. The `rawio` interface will remain throughout the life of Red Hat Enterprise Linux 5, but is a candidate for removal in a future release.

Currently, AIO (*Asynchronous I/O*) on file systems is only supported in **O_DIRECT** or non-buffered mode. Further, note that the asynchronous poll interface is no longer present, and that AIO on pipes is no longer supported.

ctmpc

ctmpc is a deprecated driver; however, it will still be included throughout the life of Red Hat Enterprise Linux 5. Note that it is a candidate for removal from future releases.

Policy Modules and semanage Support

Red Hat Enterprise Linux 5 now supports *policy modules* and *semanage*. Policy modules simplify the creation and distribution of policy customizations and third-party policies through the use of the **semodule** and **checkmodule** tools.

The **semanage** tool is a policy management tool that modifies the SELinux configuration. It also allows you to configure file contexts, networking component labeling, and user mappings for Linux-to-SELinux.

raw Device Mapping

The raw devices interface has been deprecated in Red Hat Enterprise Linux 5; raw device mapping is now configured via **udev** rules.

To configure raw device mapping, add the appropriate entries to **/etc/udev/rules.d/60-raw.rules** in the following formats:

- For device names:

```
ACTION=="add", KERNEL="<device name>", RUN+="raw /dev/raw/rawX %N"
```

- For major / minor numbers:

```
ACTION=="add", ENV{MAJOR}="A", ENV{MINOR}="B", RUN+="raw /dev/raw/rawX %M %m"
```

Replace *<device name>* with the name of the device you need to bind (for example, **/dev/sda1**). "A" and "B" are the major / minor numbers of the device you need to bind, and *X* is the raw device number that you want the system to use.

If you have a large, pre-existing **/etc/sysconfig/rawdevices** file, convert it with the following script:

```
#!/bin/sh

grep -v "^ *#" /etc/sysconfig/rawdevices | grep -v "^$" | while read dev
major minor ; do
  if [ -z "$minor" ]; then
    echo "ACTION==\"add\", KERNEL==\"${major##/dev/}\",
RUN+=\"/usr/bin/raw $dev %N\""
  else
    echo "ACTION==\"add\", ENV{MAJOR}==\"$major\", ENV{MINOR}==\"$minor\",
RUN+=\"/usr/bin/raw $dev %M %m\""
  fi
done
```

QLogic Support

Red Hat Enterprise Linux 5 supports the *QLogic* family of iSCSI HBA (*Host Bus Adapters*). At present, only the iSCSI interface to these boards are supported (using the `qla4xxx` driver).

In addition, Red Hat does not currently support these boards as Ethernet NIC, as this capability requires the `qla3xxx` driver. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.

IBM System z Instruction Set

In order to optimally exploit the *IBM System z* instruction set for 31-bit applications, it is recommended that you use the `gcc` option `-march=z900`. For 64-bit applications, the `gcc` will exploit the *IBM System z* instruction set by default.

iSeries Access for Linux

The *iSeries ODBC Driver* for Linux has been replaced by the *iSeries Access for Linux*, which can be downloaded at the following link:

<http://www.ibm.com/eserver/series/access/linux/>

The *iSeries Access for Linux* offers Linux-based access to iSeries servers, and allows you to:

- Access the DB2 UDB (*Universal Database*) for iSeries using its ODBC Driver
- Establish a 5250 session to an iSeries server from a Linux client
- Access the DB2 UDB via the EDRS (*Extended Dynamic Remote SQL*) driver
- Support 32-bit (i386 and PowerPC) and 64-bit (x86-64 and PowerPC) platforms

IBM Power4 iSeries

Red Hat Enterprise Linux no longer supports the *IBM Power4 iSeries*.

2.5. Driver Update Program

This section includes information on the implementation of the Red Hat Enterprise Linux 5 Driver Update Program.

Kernel Module Packages

On Red Hat Enterprise Linux 5, it is possible to build updated kernel module packages that depend upon the current kernel ABI version and not on a specific kernel release number. This facilitates building kernel modules that can be used against a range of Red Hat Enterprise Linux 5 kernels, rather than a single release. The project website at <http://www.kerneldrivers.org/> contains more information about the packaging process, as well as several examples.

Note that the following issues have also been identified:

- Bootpath drivers distributed as `kmod` packages are not officially supported.
- Overriding existing in-kernel drivers are not currently supported.

These issues will be addressed in a future update of Red Hat Enterprise Linux 5.

Kernel Module Loading

The module loading behavior on Red Hat Enterprise Linux 5 has changed from previous releases of Red Hat Enterprise Linux. The modules shipped in the Red Hat Enterprise Linux 5 kernel package are signed, as was the case in Red Hat Enterprise Linux 4. On Red Hat Enterprise Linux 5 kernels, however, it is no longer possible to load a signed module from another kernel build.

This means that a module shipped with the initial Red Hat Enterprise Linux 5 distribution cannot be loaded in future updated kernels. This helps prevent users from loading unsupported modules on a system. Red Hat only supports modules that are signed and included in a distribution.

If you want to load an older module, you can try rebuilding it without a signature. Alternatively, you can remove the signature from the binary file using the following command:

```
objcopy -R .module_sig <module name>-mod.ko <module name>-nosig.ko
```

It is recommended that you consult with a designated Red Hat Global Support Representative before attempting to load unsigned modules.

2.6. Internationalization

This section includes information on language support under Red Hat Enterprise Linux 5.

Input Methods

SCIM (Smart Common Input Method) has replaced IIMF as the input method system for Asian and other languages in this release. The default GTK Input Method Module for SCIM is provided by `scim-bridge`; in Qt, it is provided by `scim-qtim`.

Below are the default trigger hotkeys for different languages:

- All languages: **Ctrl-Space**
- Japanese: **Zenkaku-Hankaku** or **Alt-`**
- Korean: **Shift-Space**

If SCIM is installed, it runs by default for all users.

After installing or removing SCIM engine packages, it is recommended to start a new desktop session in order for the changes to be reflected in the SCIM language menu.

Language Installation

To enable additional language support for some Asian languages, you need to install the necessary language support packages. Below is a list of these languages and the command you need to run (as root) to install their corresponding language support packages:

- Assamese – `yum install fonts-bengali m17n-db-assamese scim-m17n`
- Bengali – `yum install fonts-bengali m17n-db-bengali scim-m17n`
- Chinese – `yum install fonts-chinese scim-chewing scim-pinyin scim-tables-chinese`
- Gujarati – `yum install fonts-gujarati m17n-db-gujarati scim-m17n`
- Hindi – `yum install fonts-hindi m17n-db-hindi scim-m17n`

- Japanese – `yum install fonts-japanese scim-anthy`
- Kannada – `yum install fonts-kannada m17n-db-kannada scim-m17n`
- Korean – `yum install fonts-korean scim-hangul`
- Malayalam – `yum install fonts-malayalam m17n-db-malayalam scim-m17n`
- Marathi – `yum install fonts-hindi m17n-db-marathi scim-m17n`
- Oriya – `yum install fonts-oriya m17n-db-oriya scim-m17n`
- Punjabi – `yum install fonts-punjabi m17n-db-punjabi scim-m17n`
- Sinhala – `yum install fonts-sinhala m17n-db-sinhala scim-m17n`
- Tamil – `yum install fonts-tamil m17n-db-tamil scim-m17n`
- Telugu – `yum install fonts-telugu m17n-db-telugu scim-m17n`

It is also recommended that you install `scim-bridge-gtk` and `scim-qtimm` when enabling additional language support. The `scim-bridge-gtk` package prevents possible binary conflicts with third-party applications linked against older versions of `libstdc++`.

Note that additional language support packs are also available for OpenOffice (`openoffice.org-langpack-<language code>_<locale>`) and KDE (`kde-i18n-<language>`). These packages can also be installed through `yum`.

im-chooser

A new user configuration tool called `im-chooser` has been added, which allows you to easily disable or enable the usage of input methods on your desktop. So if SCIM is installed but you do not wish to run it on your desktop, you can disable it using `im-chooser`.

xinputrc

At X startup, `xinput.sh` now sources `~/.xinputrc` or `/etc/X11/xinit/xinputrc` instead of searching config files under `~/.xinput.d/` or `/etc/xinit/xinput.d/`.

Pango Support in Firefox

Firefox in Red Hat Enterprise Linux 5 is built with Pango, which provides better support for certain scripts, such as Indic and some CJK scripts.

To disable the use of Pango, set `MOZ_DISABLE_PANGO=1` in your environment before launching Firefox.

Fonts

Support is now available for synthetic emboldening of fonts that do not have a bold face.

New fonts for Chinese have been added: AR PL ShanHeiSun Uni (`uming.ttf`) and AR PL ZenKai Uni (`ukai.ttf`). The default font is AR PL ShanHeiSun Uni, which contains embedded bitmaps. If you prefer outline glyphs, add the following section to your `~/.font.conf` file:

```
<fontconfig>
  <match target="font">
```



```

<test name="family" compare="eq">
  <string>AR PL ShanHeiSun Uni</string>
</test>
<edit name="embeddedbitmap" mode="assign">
  <bool>>false</bool>
</edit>
</match>
</fontconfig>

```

gtk2 IM submenu

The Gtk2 context menu IM submenu no longer appears by default. You can enable it on the command line with the following command:

```
gconftool-2 --type bool --set
'/desktop/gnome/interface/show_input_method_menu' true
```

Support for text installation on CJK

CJK (Chinese, Japanese, and Korean) rendering support has been removed from the Anaconda text installation. The text installation method is being deprecated in the long term, as the GUI installation, VNC and kickstart methods are preferred.

gtk+ deprecation

The following packages are deprecated and scheduled for removal in Red Hat Enterprise Linux:

- **gtk+**
- **gdk-pixbuf**
- **glib**

These packages are being deprecated in favor of the **gtk2** stack, which offers better functionality particularly in terms of internationalization and font handling.

CJK input on console

If you need to display Chinese, Japanese, or Korean text on the console, you need to set up a framebuffer; afterwards, install **bogl-bterm**, and run **bterm** on the framebuffer.

2.7. Kernel Notes

This section notes the differences between 2.6.9 (on which Red Hat Enterprise Linux 4 is based) and 2.6.18 (which Red Hat Enterprise Linux 5 will inherit) as of July 12, 2006. Additional features which we are currently working on upstream (for example, virtualization) that will appear late in 2.6.18 or 2.6.19 are not highlighted here. In other words, this list only shows what is already included in the upstream Linus tree; not what is currently in development. Consequently, this list is not a final, or complete list of the new Red Hat Enterprise Linux 5 features, although it does give a good overview of what can be expected. Also, note that this section only picks out highlights of upstream changes, and as such it is not fully comprehensive. It does not include mention of several low-level hardware support enhancements and device driver info.

The following is a good source for a next level-of-detail view:

<http://kernelnewbies.org/LinuxChanges>

Performance / Scalability

- Big Kernel Lock preemption (2.6.10)
- Voluntary preemption patches (2.6.13) (subset in Red Hat Enterprise Linux 4)
- Lightweight user-space priority inheritance (PI) support for futexes, useful for real-time applications (2.6.18)
 - ref: <http://lwn.net/Articles/178253/>
- New 'mutex' locking primitive (2.6.16)
- High resolution timers (2.6.16)
 - In contrast to the low-resolution timeout API implemented in kernel/timer.c, hrtimers provide finer resolution and accuracy depending on system configuration and capabilities. These timers are currently used for itimers, POSIX timers, nanosleep and precise in-kernel timing.
- Modular, on-the-fly switchable I/O schedulers (2.6.10)
 - This was adjustable only by boot option in Red Hat Enterprise Linux 4 (also system-wide instead of per-queue).
- New Pipe implementation (2.6.11)
 - 30-90% performance improvement in pipe bandwidth
 - circular buffer allows more buffering than blocking writers
- "Big Kernel Semaphore": turns the Big Kernel Lock into a semaphore
 - reduces latency by breaking up long lock hold times and adding voluntary preemption
- X86 "SMP alternatives"
 - optimizes a single kernel image at runtime according to the available platform
 - ref: <http://lwn.net/Articles/164121/>
- kernel-headers package
 - replaces the `glibc-kernheaders` package
 - provides better suitability with the new `headers_install` feature of the 2.6.18 kernel
 - notable kernel header-related changes:
 - removed `<linux/compiler.h>` header file, as it is no longer useful
 - removed `_syscallX()` macros; user-space should use `syscall()` from the C library instead
 - removed `<asm/atomic.h>` and `<asm/bitops.h>` header files; C compiler provides its own atomic built-in functions better suitable for user-space programs

- content previously protected with `#ifdef __KERNEL__` is now removed completely with the `unifdef` tool; defining `__KERNEL__` in order to view parts which should not be visible to user-space is no longer effective
- removed the `PAGE_SIZE` macro from some architectures, due to variance in page sizes; user-space should be using `sysconf (_SC_PAGE_SIZE)` or `getpagesize()`
- to provide better suitability for user-space, removed several header files and header content

Generic Feature Additions

- `kexec` and `kdump` (2.6.13)
 - `diskdump` and `netdump` have been replaced by `kexec` and `kdump`, which ensure faster boot-up and creation of reliable kernel vmcores for diagnostic purposes. For more information and configuration instructions, please refer to `/usr/share/doc/kexec-tools-<version>/kexec-kdump-howto.txt` (replace `<version>` with the corresponding version of the `kexec-tools` package installed).
 - Note that at present, virtualized kernels cannot use the `kdump` function.
- `inotify` (2.6.13)
 - user interface for this is through the following syscalls: `sys_inotify_init`, `sys_inotify_add_watch`, and `sys_inotify_rm_watch`.
- Process Events Connector (2.6.15)
 - reports `fork`, `exec`, `id` change, and `exit` events for all processes to user-space.
 - Applications that may find these events useful include accounting / auditing (for example, ELSA), system activity monitoring (for example, `top`), security, and resource management (for example, CKRM). Semantics provide the building blocks for features like per-user-namespace, "files as directories" and versioned file systems.
- Generic RTC (RealTime Clock) subsystem (2.6.17)
- `splice` (2.6.17)
 - new IO mechanism which avoids data copies when transferring data between applications
 - ref: <http://lwn.net/Articles/178199/>

File System / LVM

- EXT3
 - support for Extended Attributes in the body of large inode in ext3: saves space and improves performance in some cases (2.6.11)
- Device mapper multipath support
- ACL support for NFSv3 and NFSv4 (2.6.13)

- NFS: supports large reads and writes on the wire (2.6.16)
 - The Linux NFS client now supports transfer sizes of up to 1MB.
- VFS changes
 - The "shared subtree" patches have been merged. (2.6.15)
 - ref: <http://lwn.net/Articles/159077/>
- Big CIFS update (2.6.15)
 - features several performance improvements as well as support for Kerberos and CIFS ACL
- autofs4: updated to provide direct mount support for user-space autofs (2.6.18)
- cachefs core enablers (2.6.18)

Security

- Multilevel security implementation for SELinux (2.6.12)
- Audit subsystem
 - support for process-context based filtering (2.6.17)
 - more filter rule comparators (2.6.17)
- TCP/UDP getpeercon: enabled security-aware applications to retrieve the entire security context of a process on the other side of a socket using an IPSec security association. If only MLS-level information is needed or interoperability with legacy unix system is required, NetLabel can be used in place of IPSec.

Networking

- Added several TCP congestion modules (2.6.13)
- IPv6: supports several new sockopt / ancillary data in Advanced API (2.6.14)
- IPv4/IPv6: UFO (UDP Fragmentation Offload) Scatter-gather approach (2.6.15)
 - UFO is a feature wherein the Linux kernel network stack will offload the IP fragmentation functionality of large UDP datagram to hardware. This will reduce the overhead of stack in fragmenting the large UDP datagram to MTU-sized packets.
- Added nf_contrack subsystem (2.6.15)
 - The existing connection tracking subsystem in netfilter can only handle ipv4. There were two choices present to add connection tracking support for ipv6; either duplicate all of the ipv4 connection tracking code into an ipv6 counterpart, or (the choice taken by these patches) design a generic layer that could handle both ipv4 and ipv6 and thus requiring only one sub-protocol (TCP, UDP, etc.) connection tracking helper module to be written. In fact, nf_contrack is capable of working with any layer 3 protocol.
- IPV6
 - RFC 3484-compliant source address selection (2.6.15)

- added support for Router Preference (RFC4191) (2.6.17)
- added Router Reachability Probing (RFC4191) (2.6.17)
- added support for Multiple Routing Tables and Policy Routing
- Wireless updates
 - hardware crypto and fragmentation offload support
 - QoS (WME) support, "wireless spy support"
 - mixed PTK/GTK
 - CCMP/TKIP support and WE-19 HostAP support
 - BCM43xx wireless driver
 - ZD1211 wireless driver
 - WE-20, version 20 of the Wireless Extensions (2.6.17)
 - added the hardware-independent software MAC layer, "Soft MAC" (2.6.17)
 - added LEAP authentication type
- Added generic segmentation offload (GSO) (2.6.18)
 - can improve performance in some cases, though it needs to be enabled through `ethtool`
- DCCPv6 (2.6.16)

Added Hardware Support



NOTE

This section only enumerates the most generic features among many.

- x86-64 clustered APIC support (2.6.10)
- Infiniband support (2.6.11)
- Hot plug
 - added generic memory add/remove and supporting functions for memory hotplug (2.6.15)
- SATA/libata enhancements, additional hardware support
 - A completely reworked libata error handler; the result of all this work should be a more robust SATA subsystem which can recover from a wider range of errors.
 - Native Command Queuing (NCQ), the SATA version of tagged command queuing - the ability to have several I/O requests to the same drive outstanding at the same time. (2.6.18)

- Hotplug support (2.6.18)
- EDAC support (2.6.16)
 - The EDAC goal is to detect and report errors that occur within the system.
- Added a new ioatdma driver for the Intel(R) I/OAT DMA engine (2.6.18)

NUMA (Non-Uniform Memory Access) / Multi-core

- Cpusets (2.6.12)
 - Cpusets now provide a mechanism for assigning a set of CPUs and Memory Nodes to a set of tasks. Cpusets constrain the CPU and memory placement of tasks only to the resources within a task's current cpuset. These are essential in managing dynamic job placement on large systems.
- NUMA-aware slab allocator (2.6.14)
 - This creates slabs on multiple nodes and manages slabs in such a way that locality of allocations is optimized. Each node has its own list of partial, free and full slabs. All object allocations for a node occur from node-specific slab lists.
- Swap migration (2.6.16)
 - Swap migration allows the moving of physical location of pages between nodes in a NUMA system while the process is running.
- Huge pages (2.6.16)
 - Added NUMA policy support for huge pages: the `huge_zonelist()` function in the memory policy layer provides a list of zones ordered by NUMA distance. The `hugetlb` layer will walk that list looking for a zone that has available huge pages but is also in the nodeset of the current cpuset.
 - Huge pages now obey cpusets.
- Per-zone VM counters
 - provide zone-based VM statistics, which are necessary in determining what state of memory a zone is in
- Netfilter `ip_tables`: NUMA-aware allocation. (2.6.16)
- Multi-core
 - Added a new scheduler domain for representing multi-core with shared caches between cores. This makes it possible to make smarter cpu scheduling decisions on such systems, improving performance greatly for some cases. (2.6.17)
 - Power saving policy for the CPU scheduler: with multicore/smt cpus, the power consumption can be improved by leaving some packages idle while others do all the work, instead of spreading the tasks over all CPUs.

3. RELEASE NOTES FOR IA64

3.1. Installation-Related Notes

The following section includes information specific to the installation of Red Hat Enterprise Linux and the Anaconda installation program.



NOTE

In order to upgrade an already-installed Red Hat Enterprise Linux, you must use Red Hat Network to update those packages that have changed.

You may use Anaconda to perform a fresh installation of Red Hat Enterprise Linux 5 or to perform an upgrade from *the latest updated version* of Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5.

If you are copying the contents of the Red Hat Enterprise Linux 5 CD-ROMs (in preparation for a network-based installation, for example) be sure to copy the CD-ROMs for the operating system *only*. Do not copy the Supplementary CD-ROM, or any of the layered product CD-ROMs, as this will overwrite files necessary for Anaconda's proper operation. These CD-ROMs must be installed *after* Red Hat Enterprise Linux has been installed.

3.1.1. ISO Contents and Registration

The organization of software component packages into product-specific variants has changed from previous versions of Red Hat Enterprise Linux. The total number of different variants and ISO images has been reduced to two:

- Red Hat Enterprise Linux 5 Server
- Red Hat Enterprise Linux 5 Client

The ISO images contain software packages for a number of optional repositories that provide additional functionality over the core distribution, such as Virtualization, Clustering or Cluster Storage. For more information about the Server variants, Client variants and available options, please refer to <http://www.redhat.com/rhel/>.

Note that Virtualization is currently being offered in this release as a technology preview.

With optional content in the same tree or ISO image, it is important to avoid a mismatch between the components offered for installation and those covered by the subscription. Such a mismatch could result in an increased exposure to bug and vulnerability risks.

In order to ensure that the components offered are in sync with the subscription, Red Hat Enterprise Linux 5 requires entering an *Installation Number* that will be used to configure the installer to offer the right package set. This Installation Number is included in your subscription.

If you skip entering the Installation Number, this will result in a core Server or Desktop installation. Additional functionality can then be added manually at a later time. For more information about Installation Numbers, please refer to <http://www.redhat.com/apps/support/in.html>.

The Installation Number used during the installation process will be saved in `/etc/sysconfig/rhn/install-num`. When registering with Red Hat Network, this file will be referenced by `rhn_register` to automatically determine which appropriate child channels the system should be subscribed to.

3.1.2. New RPM GPG Signing Keys

A new release signing key is used to sign Red Hat Enterprise Linux 5 packages. When updating a system for the first time, you will be prompted to allow this key to be installed.

Signing keys are distributed in the following files:

- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release` – contains the public key for the new release signing key
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-auxiliary` – contains the public key for an auxiliary release signing key, currently not in use
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-former` – contains the public key for the previous release signing key, used for past Red Hat Enterprise Linux releases

3.1.3. Subversion

In Red Hat Enterprise Linux 5, the Subversion version control system is linked against Berkeley DB 4.3. If you are upgrading from Red Hat Enterprise Linux 4 and any Subversion repositories which use the Berkeley DB backend "BDB" (rather than the pure file system-based "FSFS" backend) have been created on the system, special care must be taken to ensure the repositories can be accessible after the upgrade. The following process must be performed on the Red Hat Enterprise Linux 4 system, *prior* to upgrading to Red Hat Enterprise Linux 5:

1. Shut down any running processes and ensure that no processes can access the repository (for example, `httpd`, `svnserve` or any local users with direct access).
2. Create a backup of the repository using the following command:

```
svnadmin dump /path/to/repository | gzip > repository-backup.gz
```

3. Run the `svnadmin recover` command on the repository:

```
svnadmin recover /path/to/repository
```

4. Delete any unused log files in the repository:

```
svnadmin list-unused-dblogs /path/to/repository | xargs rm -vf
```

5. Delete any remaining shared-memory files in the repository:

```
rm -f /path/to/repository/db/___db.0*
```

3.1.4. Other Installation Notes

- If IDE/PATA (*Parallel ATA*) devices are configured in "100% Native" mode, some BIOSes may prevent the Red Hat Enterprise Linux 5 installation process from completing successfully. To prevent this from occurring, configure the IDE/PATA mode as "Legacy" in the BIOS.
- The *IBM System z* does not provide a traditional Unix-style physical console. As such, Red Hat Enterprise Linux 5 for the *IBM System z* does not support the `firstboot` functionality during initial program load.

To properly initialize setup for Red Hat Enterprise Linux 5 on the *IBM System z*, run the following commands after installation:

- `/usr/bin/setup` – provided by the `setuptools` package
- `/usr/bin/rhn_register` – provided by the `rhn-setup` package
- Red Hat Enterprise Linux 5 for the 64-bit Intel Itanium2 architecture includes runtime support for 32-bit applications through the use of Intel's IA-32 Execution Layer.

The IA-32 Execution Layer is provided on the Supplementary disc for the Intel Itanium2 architecture. In addition, a set of 32-bit libraries and applications are provided on a separate 32-bit Compatibility Layer disc. The IA-32 Execution Layer and 32-bit compatibility packages together provide a runtime environment for 32-bit applications on the 64-bit native distribution.

To install the IA-32 Execution Layer and required 32-bit compatibility packages, follow these steps:

1. Install Red Hat Enterprise Linux 5 for the Intel Itanium2 Architecture.
 2. Insert the Red Hat Enterprise Linux 5 Supplementary CD, which contains the `ia32e1` package.
 3. After the system has mounted the CD, change to the directory containing the **Supplementary packages**. For example:


```
cd /media/cdrom/Supplementary/
```
 4. Install the `ia32e1` package:


```
rpm -Uvh ia32e1-<version>.ia64.rpm
```

Replace `<version>` with the corresponding version of the `ia32e1` package to be installed.
 5. Eject the Supplementary CD:


```
eject /media/cdrom
```
 6. To verify the installation of the 32-bit compatibility layer and libraries after installation, check that the `/emu1` directory has been created and that it contains files.
 7. To verify that the 32-bit compatibility mode is in effect, type the following in a shell prompt:


```
service ia32e1 status
```
 8. At this point you can install compatibility libraries by inserting the 32-bit Compatibility Layer disc. You may choose to install all of the packages available on the disc or choose the particular packages required in order to provide runtime support for your 32-bit applications.
- When booting Anaconda with PXE using the parameter `ksdevice=bootif`, you will still be prompted for the ethernet interface to use during installation. If only one ethernet device is plugged in, use the parameter `ksdevice=link` instead. Alternatively, you can also specify the interface manually.

3.2. Technology Previews

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 5 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. Erratas 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.

Stateless Linux

Included in this release of Red Hat Enterprise Linux 5 are enabling infrastructure pieces for *Stateless Linux*. Stateless Linux is a new way of thinking about how a system is to be run and managed, designed to simplify provisioning and management of large numbers of systems by making them easily replaceable. This is accomplished primarily by establishing prepared system images which get replicated and managed across a large number of stateless systems, running the operating system in a read-only manner (please refer to `/etc/sysconfig/readonly-root` for more details).

In its current state of development, the Stateless features are subsets of the intended goals. As such, the capability is being labeled as a technology preview.

The following is a list of the initial capabilities included in Red Hat Enterprise Linux 5:

- running a stateless image over NFS
- running a stateless image via loopback over NFS
- running on iSCSI

It is highly recommended that those interested in testing stateless code read the HOWTO at <http://fedoraproject.org/wiki/StatelessLinuxHOWTO> and join stateless-list@redhat.com.

GFS2

GFS2 is an evolutionary advancement based on the GFS file system. While fully functional, GFS2 is not yet considered production-ready. GFS2 is targeted to become fully supported in a subsequent Red Hat Enterprise Linux 5 update. There is also an in-place conversion utility, `gfs2_convert`, which can update the metadata of the older GFS file system format, converting it to a GFS2 file system.

FS-Cache

FS-Cache is a local caching facility for remote file systems that allows users to cache NFS data on a locally mounted disk. To set up the FS-Cache facility, install the `cachefilesd` RPM and refer to the instructions in `/usr/share/doc/cachefilesd-<version>/README`.

Replace `<version>` with the corresponding version of the `cachefilesd` package installed.

Virtualization

Red Hat Enterprise Linux 5 features a technology preview of Xen-based virtualization capabilities for IA64, as well as the software infrastructure needed to manage a virtualized environment.

The implementation of Virtualization in Red Hat Enterprise Linux 5 is based on the *hypervisor*, which

facilitates extremely low overhead virtualization through paravirtualization. With Intel VT-I capable processors, virtualization in Red Hat Enterprise Linux 5 also allows operating systems to run unmodified in fully virtualized mode.

Virtualization on Red Hat Enterprise Linux 5 also features the following:

- *Libvirt*, a library that provides a consistent, portable API for managing virtual machines.
- *Virtual Machine Manager*, a graphical utility for monitoring and managing virtual machines.
- Virtual machine setup through the installer, as well as the ability to kickstart virtual machines.

To install the Virtualization feature, you need to use the `linux debug` parameter when initializing Anaconda.

Virtual machines can also be managed through Red Hat Network.

At present, the Virtualization feature has the following limitations:

- When Virtualization is enabled, neither suspend to RAM nor suspend to disk are supported, and CPU frequency scaling cannot be performed.
- Fully virtualized guests cannot be saved, restored or migrated.
- The `xm create` command does not have a graphical equivalent in Virtual Machine Manager.
- Virtualization only supports the *bridged networking* component. All corresponding tools used by guests automatically choose this as the default.
- The default Red Hat SELinux policy for Virtualization only allows configuration files to be written to `/etc/xen`, log files to be written to `/var/log/xen/`, and disk files (including core dumps) to be written to `/var/lib/xen`. These defaults can be changed using the `semanage` tool.
- The hypervisor included in this release of Virtualization is *not* NUMA-aware; as such, its performance on NUMA machines may be sub-optimal. This will be addressed in a future update of Red Hat Enterprise Linux 5.

To work around this, enable `memory node interleaving` in the NUMA machine's BIOS. This ensures a more consistent performance.

- Paravirtualized domains currently do not support keymaps other than `en-US`. As such, other keyboards may not be able to type certain keystrokes. This will be addressed in a future update of Red Hat Enterprise Linux 5.
- The virtualized kernel cannot use the `kdump` function.
- `qcow` and `vmdk` images are not supported. When manually configuring guests, images backed by a physical or logical device should use the `phy:` type. For file-backed images, set the image type to `tap:aio:` for paravirtualized guests and `file:` for fully virtualized guests.
- Installing the `kernel-xen` package may lead to incorrect or incomplete `elilo.conf` entries. As such, you may need to manually edit `elilo.conf` before booting the installed `kernel-xen` package.

For example, after installing the `kernel-xen` package, `elilo.conf` may read as follows:

```
image=vmlinuz-2.6.18-1.2732.el5xen
    label=linux
    initrd=initrd-2.6.18-1.2732.el5xen.img
    read-only
    root=/dev/VolGroup01/LogVol00
    append="rhgb quiet"
```

In this instance, the following line is missing, and will need to be added prior to boot:

```
vmm=xen.gz-2.6.18-1.2747.el5
```

- Paravirtualized domains can only auto-detect relative mouse movement, and pointer movement is rather erratic. This will be addressed in a future update of Red Hat Enterprise Linux 5.
- In order to have a working console for a paravirtualized guest, you need to specify `console=xvc0` in the kernel command line.
- When guest operating systems are configured to use sparse files, `dom0` can run out of disk space. Such occurrences prevent guest disk writes from completing, and can cause data loss in guests. Further, guests that use sparse files do not synchronize I/O safely.

As such, it is recommended that you use non-sparse files instead. To configure guests to use non-sparse files, use the option `--nonsparse` when conducting a `virt-install`.

Compiz

Compiz is an OpenGL-based compositing window manager. In addition to regular window management, `compiz` also acts as a compositing manager, coordinating and synchronizing the overall desktop redrawing to provide a smoother desktop experience with less flicker.

`Compiz` uses 3D hardware acceleration to render effects such as live thumbnail windows, window drop shadows, animated window minimizing and transitions between virtual desktops.

Due to limitations in the current rendering architecture, `compiz` cannot work correctly with direct rendering OpenGL applications or applications using the Xv extension. Such applications will exhibit harmless rendering artifacts; because of this, `compiz` is currently a technology preview.

Enhancement for Ext3

In Red Hat Enterprise Linux 5, the EXT3 file system capacity has been extended beyond 8TB to a maximum of 16TB. This capability is being included as a technology preview, and is targeted for full support in a future release of Red Hat Enterprise Linux 5.

AIGLX

AIGLX is a technology preview feature of the otherwise fully supported X server. It aims to enable GL-accelerated effects on a standard desktop. The project consists of the following:

- a lightly modified X server
- an updated Mesa package that adds new protocol support

By installing these components, you can have GL-accelerated effects on your desktop with very few changes, as well as the ability to enable and disable them at will without replacing your X server. AIGLX also enables remote GLX applications to take advantage of hardware GLX acceleration.

Frysk GUI

The goal of the *frysk* project is to create an intelligent, distributed, always-on system monitoring and debugging tool that allows developers and system administrators to:

- monitor running processes and threads (including creation and destruction events)
- monitor the use of locking primitives
- expose deadlocks
- gather data
- debug any given process by choosing it from a list or allowing *frysk* to open a source code (or other) window on a process that is crashing or misbehaving

In Red Hat Enterprise Linux 5 the *frysk* graphical user interface is a technology preview, whereas the *frysk* command line interface is fully supported.

Systemtap

Systemtap provides free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. With the help of *systemtap*, developers no longer need to go through the tedious and disruptive instrument, recompile, install, and reboot sequence that may be otherwise required to collect data.

Dogtail

Dogtail is a GUI test tool and automation framework written in Python that uses *Accessibility* technologies to communicate with desktop applications.

Support for Indic Languages and Sinhalese

Red Hat Enterprise Linux 5 also features support for the following languages as technology preview:

- *Assamese*
- *Kannada*
- *Sinhalese*
- *Telugu*

For more information about how to install and enable support for these languages, refer to the *Internationalization* section of this document.

Installing to dm-multipath Devices

Anaconda now has the capability to detect, create, and install to *dm-multipath* devices. To enable this feature, add the parameter `mpath` to the kernel boot line.

Note that the parameter `mpath` may cause a boot failure if a device's `major:minor` number changes. This issue will be addressed in a future update of Red Hat Enterprise Linux 5.

Installation / Boot for iSCSI software initiator (open-iscsi)

Anaconda now provides the ability to install to an iSCSI device. Booting and installing is fully supported with the *QLogic qla4xxx* hardware initiator. However, the capability to install to an iSCSI device for the `open-iscsi` software initiator is currently considered a Technology Preview, due to the following issues:

- Text mode installation does not complete. You must do a graphical install, or an automated kickstart install.
- Media-based installations do not complete. You must do a network-based install.
- Depending on the timing of events, Anaconda may be unable to detect all the iSCSI targets or LUNs. When this occurs, use the installer shell to configure the storage through iSCSI commands.
- The `iscsid` daemon may not properly start. Such an occurrence will prevent the system from handling all iSCSI errors, such as network problems, SCSI/iSCSI timeouts, and target errors. To confirm that the `iscsid` daemon is running, run the command `iscsiadm -m session -i` and check that the line, **Internal iscsid Session State:** reports a value (it can be any value).
- On certain iSCSI target implementations, the system may hang during shutdown.
- On certain iSCSI target implementations, the system may hang during reboot. To avoid this, shutdown the system and boot it up again (instead of rebooting directly from a session).
- Booting from iSCSI devices on the *IBM System p* does not work reliably. While installation on an iSCSI device may appear to succeed, the resulting installation will not boot properly.
- On the first boot after install, you may receive SELinux errors such as the following:

```
kernel: audit(1169664832.270:4): avc: denied { read
} for pid=1964 comm="iscsid"
```

To work around this, boot the system with the kernel parameter `enforcing=0`. Once the system has properly booted, use the command `setenforce 1` to restore enforcing mode.

These limitations will be addressed in a future Red Hat Enterprise Linux 5 update.

3.3. Known Issues

- Host bus adapters that use the MegaRAID driver must be set to operate in "Mass Storage" emulation mode, not in "I2O" emulation mode. To do this, perform the following steps:
 1. Enter the MegaRAID BIOS Set Up Utility.
 2. Enter the Adapter settings menu.
 3. Under Other Adapter Options, select Emulation and set it to Mass Storage.

If the adapter is incorrectly set to "I2O" emulation, the system will attempt to load the i2o driver. This will fail, and prevent the proper driver from being loaded.

Previous Red Hat Enterprise Linux releases generally do not attempt to load the I2O driver before the MegaRAID driver. Regardless of this, the hardware should always be set to "Mass Storage" emulation mode when used with Linux.

- Red Hat Enterprise Linux 5 includes `openmpi-1.1.1-4.el5` (from the OFED 1.1 distribution), which has been discovered to eventually quit working entirely. This happens after the `openmpi` stack works as expected for a varying amount of time.

For updated versions of `openmpi`, please check <http://people.redhat.com/dledford/Infiniband/openmpi>

- Installing Windows Server 2003 as a guest on a fully virtualized Red Hat Enterprise Linux 5 system ends unexpectedly after completing the first stage of installation. When this happens, the graphical console window closes, and the guest disappears from the **Virtual Machine Manager's** list of machines, resulting in a **Broken pipe** error.

This issue will be resolved in an upcoming Red Hat Enterprise Linux 5 update. To work around this, use the following command at the terminal:

```
xm create /etc/xen/<name of guest machine>
```

Afterwards, open the virtual machine.

- When attempting to create a fully virtualized Windows Server 2003 from a CD / DVD, the second stage of the guest install will not continue upon reboot.

To work around this, edit `/etc/xen/<name of guest machine>` by properly appending an entry for the CD / DVD device.

If an installation to a simple file is used as a virtual device, the `disk` line of `/etc/xen/<name of guest machine>` will read like the following:

```
disk = [ 'file:/PATH-OF-SIMPLE-FILE,hda,w' ]
```

A DVD-ROM device located on the host as `/dev/dvd` can be made available to stage 2 of the installation as `hdc` by appending an entry like `'phy:/dev/dvd,hdc:cdrom,r'`. As such, the disk line should now read as follows:

```
disk = [ 'file:/opt/win2003-sp1-20061107,hda,w',
        'phy:/dev/dvd,hdc:cdrom,r' ]
```

The precise device path to use may vary depending on your hardware.

- `rmmmod xennet` causes `domU` to crash; this is caused by a grant table issue in the Virtualization feature. Due to the current inability of the Virtualization feature to asynchronously release grant table operations, it is unsafe to unload the `xennet` module in guests. In such situations, grant tables are used to perform backend-frontend communication, and there is no guarantee that the backend will release the references, leading to an inevitable memory leak.

This issue will be resolved in the next minor release of Red Hat Enterprise Linux 5. At present, users are advised *not* to unload the `xennet` module in guests.

- Running `ethtool eth0` outputs incomplete information about the ethernet card settings. This only occurs in systems running a virtualized kernel, since the Virtualization feature uses a networking setup where the physical ethernet device is identified as `peth0`. As such, the

correct command for retrieving information about the physical ethernet device is `ethtool peth0`.

- When using Red Hat Enterprise Linux 5 on a machine with an *nVidia CK804* chipset installed, you may receive kernel messages similar to the following:

```
kernel: assign_interrupt_mode Found MSI capability
kernel: pcie_portdrv_probe->Dev[005d:10de] has invalid IRQ. Check
vendor BIOS
```

These messages indicate that certain PCI-E ports are not requesting IRQs. Further, these messages do not, in any way, affect the operation of the machine.

- Some *Cisco Aironet Wireless* devices prevent **NetworkManager** from storing connection details for wireless networks that do not broadcast an SSID. This is caused by a *Cisco Aironet Wireless* device firmware limitation.
- Laptops that have the *Cisco Aironet MPI-350* wireless card equipped may hang trying to get a DHCP address during any network-based installation using the wired ethernet port.

To work around this, use local media for your installation. Alternatively, you can disable the wireless card in the laptop BIOS prior to installation (you can re-enable the wireless card after completing the installation).

- Currently, `system-config-kickstart` does not support package selection and deselection. When using `system-config-kickstart`, the **Package Selection** option indicates that it is disabled. This is because `system-config-kickstart` uses `yum` to gather group information, but is unable to configure `yum` to connect to Red Hat Network.

This issue is currently being investigated for resolution by the next minor release of Red Hat Enterprise Linux 5. At present, you need to update package sections in your kickstart files manually. When using `system-config-kickstart` to open a kickstart file, it will preserve all package information in it and write it back out when you save.

- 4-socket AMD *Sun Blade X8400 Server Modules* systems that do not have memory configured in `node 0` will panic during boot. Systems should be configured with memory in `node 0` to prevent the kernel panic.
- Installing to LVM mirror devices through Anaconda is currently not supported. This capability will be added in a future update of Red Hat Enterprise Linux 5.
- When installing Red Hat Enterprise Linux 5 from a directory on an NFS server containing Red Hat Enterprise Linux ISO images, Anaconda may display the following error message:

```
Unable to read package metadata. This may be due to a missing
repdata directory.
Please ensure that your install tree has been correctly generated.
Cannot
open/read repomd.xml file for repository:
```

This problem occurs if the directory holding the ISO images also contains a partially unpacked installation tree (for example, the `/images` directory from the first ISO). The presence of such directories results in the error stated above.

To prevent this error, unpack trees only to directories other than the one containing the installation ISO images.

- Boot-time logging to `/var/log/boot.log` is not available in this release of Red Hat Enterprise Linux 5. An equivalent functionality will be added in a future update of Red Hat Enterprise Linux 5.
- Neither `kexec` nor `kdump` are able to dump onto disks attached to an `accraid` controller.

To work around this issue, use `scp` for network dumping. Alternatively, you can also dump onto a disk through a different controller.

- Installing a fully virtualized guest using split installation media -- specifically, multiple CD-ROMs -- may fail when required to switch between installation CDs. During the guest OS installation process, users may be prevented from mounting or ejecting installation CDs, which prevents the installation from completing.

As such, it is recommended that you use the QEMU monitor console to switch CD-ROM images during the guest OS installation process. The procedure is as follows:

1. Open a graphical VNC console to the guest OS.
2. Unmount the CD-ROM device in the guest OS.
3. Switch to the QEMU monitor console by pressing `Ctrl-Alt-2`.
4. Run the command `eject hdc`.
5. Run the command `change hdc <path to the CD-ROM in host system>`.
6. Switch back to the guest OS console by pressing `Ctrl-Alt-1`.
7. Mount the CD-ROM device in the guest OS.

Note that when using a regular VNC client the host X server may encounter some difficulty interpreting the `Ctrl-Alt-2` and `Ctrl-Alt-1` command. To work around this in `virt-manager`, use *sticky keys*. Pressing `Ctrl` three times makes it "sticky" until the next non-modifier is pressed. As such, to send `Ctrl-Alt-1`, press `Ctrl` twice before pressing `Ctrl-Alt-1`.

- The Red Hat Enterprise Linux 5 *Driver Update Model* creates modified `initrd` images whenever a `kmod` package that includes a bootpath-modifying driver is installed. In time, the number of backup `initrd` images may soon fill the `/boot` partition, particularly if the system undergoes a sizeable number of driver updates.

As such, it is recommended that you monitor the free space on the `/boot` partition if you regularly perform driver updates. You can free up more space in `/boot` by removing older `initrd` images; these files end in `.img0`, `.img1`, `.img2`, and so on.

- Autorun on removable media is currently disabled. To install packages from the Red Hat Enterprise Linux Supplementary CD, launch the CD installer manually using the following command:

```
system-cdinstall-helper /media/path-to-mounted-drive
```

- When upgrading from Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5, the Deployment Guide is not automatically installed. You need to use `pirut` to manually install it after completing the upgrade.
- An `autofs` bug prevents multi-mounts from working properly.

During an expiry, if the last multi-mount component to be checked does not have a mount associated with it while other components are busy, `autofs` erroneously determines the multi-mount to be expirable. This causes the multi-mount to be partially expired, resulting in the multi-mount becoming unresponsive to further mount requests and expire runs.

To permanently resolve this problem, update `autofs` using the command `yum update autofs`.

- The system may not successfully reboot into a `kexec/kdump` kernel if X is running and using a driver other than `vesa`. This problem only exists with *ATI Rage XL* graphics chipsets.

If X is running on a system equipped with *ATI Rage XL*, ensure that it is using the `vesa` driver in order to successfully reboot into a `kexec/kdump` kernel.

- Creating a fully virtualized guest using a `boot.iso` on an NFS share mounted as read-write will not complete correctly. To work around this problem, mount the NFS share as read-only.

If you are unable to mount the NFS share as read-only, copy the `boot.iso` to the local `/var/lib/xen/images/` directory.

3.4. General Information

This section contains general information not specific to any other section of this document.

Red Hat Enterprise Linux Deployment Guide

This release of Red Hat Enterprise Linux includes a fairly comprehensive Deployment Guide. To access it, go to **System** (on the top panel) => **Documentation** => **Red Hat Enterprise Linux Deployment Guide**.

It is the intention of Red Hat to provide fully localized versions of the Deployment Guide for all supported languages. If you have installed a localized version of the Deployment Guide, it is recommended that you update it when a new version becomes available through Red Hat Network.

Web Server Packaging Changes

Red Hat Enterprise Linux 5 now includes version 2.2 of the Apache HTTP Server. This release brings a number of improvements over the 2.0 series, including:

- improved caching modules (`mod_cache`, `mod_disk_cache`, `mod_mem_cache`)
- a new structure for authentication and authorization support, replacing the authentication modules provided in previous versions
- support for proxy load balancing (`mod_proxy_balancer`)
- support for handling large files (namely, greater than 2GB) on 32-bit platforms

The following changes have been made to the default `httpd` configuration:

- The `mod_cern_meta` and `mod_asis` modules are no longer loaded by default.

- The `mod_ext_filter` module is now loaded by default.

If you are upgrading from a previous release of Red Hat Enterprise Linux, the `httpd` configuration will need to be updated for `httpd 2.2`. For more information, refer to <http://httpd.apache.org/docs/2.2/upgrading.html>.

Note that any third-party modules compiled for `httpd 2.0` must be rebuilt for `httpd 2.2`.

php

Version 5.1 of PHP is now included in Red Hat Enterprise Linux 5, which includes a number of changes to the language along with significant performance improvements. Some scripts might need to be edited for use with the new version; please refer to the link below for more information on migrating from PHP 4.3 to PHP 5.1:

<http://www.php.net/manual/en/migration5.php>

The `/usr/bin/php` executable is now built using the CLI command-line SAPI, rather than the CGI SAPI. Use `/usr/bin/php-cgi` for CGI SAPI. The `php-cgi` executable also includes FastCGI support.

The following extension modules have been added:

- the `mysqli` extension, a new interface designed specifically for MySQL 4.1 (included in the `php-mysql` package)
- `date`, `hash`, `Reflection`, `SPL` and `SimpleXML` (built-in with the `php` package)
- `pdo` and `pdo_sqlite` (in the `php-pdo` package)
- `pdo_mysql` (in the `php-mysql` package)
- `pdo_pgsql` (in the `php-pgsql` package)
- `pdo_odbc` (in the `php-odbc` package)
- `soap` (in the `php-soap` package)
- `xmlreader` and `xmlwriter` (in the `php-xml` package)
- `dom` (replacing the `domxml` extension in the `php-xml` package)

The following extension modules are no longer included:

- `dbx`
- `dio`
- `yp`
- `overload`
- `domxml`

The PEAR Framework

The PEAR framework is now packaged in the `php-pear` package. Only the following PEAR components are included in Red Hat Enterprise Linux 5:

- `Archive_Tar`
- `Console_Getopt`
- `XML_RPC`

Encrypted Swap Partitions and Non-root File Systems

Red Hat Enterprise Linux 5 now provides basic support for encrypted swap partitions and non-root file systems. To use these features, add the appropriate entries to `/etc/crypttab` and reference the created devices in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry:

```
my_swap /dev/hdb1 /dev/urandom swap,cipher=aes-cbc-essiv:sha256
```

This creates the encrypted block device `/dev/mapper/my_swap`, which can be referenced in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry for a file system volume:

```
my_volume /dev/hda5 /etc/volume_key cipher=aes-cbc-essiv:sha256
```

The `/etc/volume_key` file contains a plaintext encryption key. You can also specify `none` as the key file name; this configures the system to ask for the encryption key during boot instead.

It is recommended to use LUKS (Linux Unified Key Setup) for setting up file system volumes. To do this, follow these steps:

1. Create the encrypted volume using `cryptsetup luksFormat`.
2. Add the necessary entry to `/etc/crypttab`.
3. Set up the volume manually using `cryptsetup luksOpen` (or reboot).
4. Create a file system on the encrypted volume.
5. Add the necessary entry to `/etc/fstab`.

mount and umount

The `mount` and `umount` commands no longer directly support NFS; a built-in NFS client no longer exists. A separate `nfs-utils` package, which provides `/sbin/mount.nfs` and `/sbin/umount.nfs` helpers, must be installed for this.

CUPS Printer Browsing

CUPS printer browsing over a local subnet can be configured using the graphical tool `system-config-printer`. It can also be done using the CUPS web interface, <http://localhost:631/>.

To use directed broadcasts for printer browsing between subnets, open `/etc/cups/cupsd.conf` on the clients and replace `BrowseAllow @LOCAL` with `BrowseAllow ALL`.

ATI and R500 Support

ATI graphics cards based on the R500 chipset are supported for the `vesa` driver only, and are not supported by Red Hat Enterprise Linux 5 on external monitors, LCD projectors or accelerated 3D support.

up2date and yum

`up2date` is being deprecated in favor of `yum` (Yellowdog Updater Modified). As such, it is advisable that you revise any `up2date`-dependent scripts your system is using accordingly. For more information about `yum`, consult its man page with the command `man yum`; you can also consult the installed documentation under the directories `/usr/share/doc/yum-<version>` and `/usr/share/doc/yum-metadata-parser-<version>` (replace `<version>` with the corresponding version of `yum` and `yum-metadata-parser` installed).

OpenLDAP Server and Red Hat Directory Server

Red Hat Directory Server is an LDAP-based server that centralizes enterprise and network data into an OS-independent, network-based registry. It is set to replace OpenLDAP server components, which will be deprecated *after* Red Hat Enterprise Linux 5. For more information about Red Hat Directory Server, refer to <http://www.redhat.com/software/rha/directory/>.

i810 Driver and i830 Support

The *i810* driver supports all integrated Intel graphics chipsets, from i810 to i965. However, the support for i830 (and newer) chipsets is limited; the i810 driver can only set modes listed in the video BIOS. If your machine has an i830 or newer chipset installed, run the following command to determine what the available modes are:

```
grep Mode: /var/log/Xorg.0.log
```

Modes marked with an asterisk (*) are available for selection.

Many laptop video BIOSes do not supply a mode that matches the native panel size. Therefore the chosen mode may appear stretched, distorted, or with black borders. As such, if your chosen mode does not display properly, you need a BIOS update from your hardware vendor for the native panel size to work correctly.

Intel PRO/Wireless 3945ABG Network Connection Support

This release of Red Hat Enterprise Linux 5 includes support for the *ipw3945* (*Intel PRO/Wireless 3945ABG Network Connection*) adapter. The Red Hat Enterprise Linux 5 Supplementary disc contains the driver, regulatory daemon and firmware needed to support this adapter.

To enable support for the *ipw3945* wireless adapter, search the Red Hat Enterprise Linux 5 Supplementary disc for packages with filenames containing "3945" and install them.

rawio

`rawio` is a deprecated interface; however, Red Hat Enterprise Linux 5 still includes support for it. If you have an application that performs device access using `rawio`, it is highly recommended that you modify your application to open the block device with the `O_DIRECT` flag. The `rawio` interface will remain throughout the life of Red Hat Enterprise Linux 5, but is a candidate for removal in a future release.

Currently, AIO (*Asynchronous I/O*) on file systems is only supported in `O_DIRECT` or non-buffered mode. Further, note that the asynchronous poll interface is no longer present, and that AIO on pipes is no longer supported.

ctmpc

ctmpc is a deprecated driver; however, it will still be included throughout the life of Red Hat Enterprise Linux 5. Note that it is a candidate for removal from future releases.

Policy Modules and semanage Support

Red Hat Enterprise Linux 5 now supports *policy modules* and *semanage*. Policy modules simplify the creation and distribution of policy customizations and third-party policies through the use of the `semodule` and `checkmodule` tools.

The `semanage` tool is a policy management tool that modifies the SELinux configuration. It also allows you to configure file contexts, networking component labeling, and user mappings for Linux-to-SELinux.

raw Device Mapping

The `raw devices` interface has been deprecated in Red Hat Enterprise Linux 5; `raw device mapping` is now configured via `udev` rules.

To configure `raw device mapping`, add the appropriate entries to `/etc/udev/rules.d/60-raw.rules` in the following formats:

- For device names:

```
ACTION=="add", KERNEL="<device name>", RUN+="raw /dev/raw/rawX %N"
```

- For major / minor numbers:

```
ACTION=="add", ENV{MAJOR}="A", ENV{MINOR}="B", RUN+="raw /dev/raw/rawX %M %m"
```

Replace `<device name>` with the name of the device you need to bind (for example, `/dev/sda1`). "A" and "B" are the major / minor numbers of the device you need to bind, and `X` is the raw device number that you want the system to use.

If you have a large, pre-existing `/etc/sysconfig/rawdevices` file, convert it with the following script:

```
#!/bin/sh

grep -v "^ *#" /etc/sysconfig/rawdevices | grep -v "^$" | while read dev
major minor ; do
  if [ -z "$minor" ]; then
    echo "ACTION==\"add\", KERNEL==\"${major##/dev/}\",
RUN+=\"/usr/bin/raw $dev %N\""
  else
    echo "ACTION==\"add\", ENV{MAJOR}==\"$major\", ENV{MINOR}==\"$minor\",
RUN+=\"/usr/bin/raw $dev %M %m\""
  fi
done
```

QLogic Support

Red Hat Enterprise Linux 5 supports the *QLogic* family of iSCSI HBA (*Host Bus Adapters*). At present, only the iSCSI interface to these boards are supported (using the `qla4xxx` driver).

In addition, Red Hat does not currently support these boards as Ethernet NIC, as this capability requires the `qla3xxx` driver. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.

IBM System z Instruction Set

In order to optimally exploit the *IBM System z* instruction set for 31-bit applications, it is recommended that you use the `gcc` option `-march=z900`. For 64-bit applications, the `gcc` will exploit the *IBM System z* instruction set by default.

iSeries Access for Linux

The *iSeries ODBC Driver* for Linux has been replaced by the *iSeries Access for Linux*, which can be downloaded at the following link:

<http://www.ibm.com/eserver/series/access/linux/>

The *iSeries Access for Linux* offers Linux-based access to iSeries servers, and allows you to:

- Access the DB2 UDB (*Universal Database*) for iSeries using its ODBC Driver
- Establish a 5250 session to an iSeries server from a Linux client
- Access the DB2 UDB via the EDRS (*Extended Dynamic Remote SQL*) driver
- Support 32-bit (i386 and PowerPC) and 64-bit (x86-64 and PowerPC) platforms

IBM Power4 iSeries

Red Hat Enterprise Linux no longer supports the *IBM Power4 iSeries*.

kdump Minimum Memory

`kdump` reserves a significant amount of memory at boot time, which changes the actual minimum memory requirements of Red Hat Enterprise Linux 5. To compute the actual minimum memory requirements for your system, refer to <http://www.redhat.com/rhel/details/limits/> for the listed minimum memory requirements; then, add the amount of memory used by `kdump` to determine the actual minimum memory requirements.

3.5. Driver Update Program

This section includes information on the implementation of the Red Hat Enterprise Linux 5 Driver Update Program.

Kernel Module Packages

On Red Hat Enterprise Linux 5, it is possible to build updated kernel module packages that depend upon the current kernel ABI version and not on a specific kernel release number. This facilitates building kernel modules that can be used against a range of Red Hat Enterprise Linux 5 kernels, rather than a single release. The project website at <http://www.kerneldrivers.org/> contains more information about the packaging process, as well as several examples.

Note that the following issues have also been identified:

- Bootpath drivers distributed as `kmod` packages are not officially supported.
- Overriding existing in-kernel drivers are not currently supported.

These issues will be addressed in a future update of Red Hat Enterprise Linux 5.

Kernel Module Loading

The module loading behavior on Red Hat Enterprise Linux 5 has changed from previous releases of Red Hat Enterprise Linux. The modules shipped in the Red Hat Enterprise Linux 5 kernel package are signed, as was the case in Red Hat Enterprise Linux 4. On Red Hat Enterprise Linux 5 kernels, however, it is no longer possible to load a signed module from another kernel build.

This means that a module shipped with the initial Red Hat Enterprise Linux 5 distribution cannot be loaded in future updated kernels. This helps prevent users from loading unsupported modules on a system. Red Hat only supports modules that are signed and included in a distribution.

If you want to load an older module, you can try rebuilding it without a signature. Alternatively, you can remove the signature from the binary file using the following command:

```
objcopy -R .module_sig <module name>-mod.ko <module name>-nosig.ko
```

It is recommended that you consult with a designated Red Hat Global Support Representative before attempting to load unsigned modules.

3.6. Internationalization

This section includes information on language support under Red Hat Enterprise Linux 5.

Input Methods

SCIM (Smart Common Input Method) has replaced IIMF as the input method system for Asian and other languages in this release. The default GTK Input Method Module for SCIM is provided by `scim-bridge`; in Qt, it is provided by `scim-qtim`.

Below are the default trigger hotkeys for different languages:

- All languages: **Ctrl-Space**
- Japanese: **Zenkaku-Hankaku** or **Alt-`**
- Korean: **Shift-Space**

If SCIM is installed, it runs by default for all users.

After installing or removing SCIM engine packages, it is recommended to start a new desktop session in order for the changes to be reflected in the SCIM language menu.

Language Installation

To enable additional language support for some Asian languages, you need to install the necessary language support packages. Below is a list of these languages and the command you need to run (as root) to install their corresponding language support packages:

- Assamese – `yum install fonts-bengali m17n-db-assamese scim-m17n`
- Bengali – `yum install fonts-bengali m17n-db-bengali scim-m17n`
- Chinese – `yum install fonts-chinese scim-chewing scim-pinyin scim-tables-chinese`

- `Gujarati-yum install fonts-gujarati m17n-db-gujarati scim-m17n`
- `Hindi-yum install fonts-hindi m17n-db-hindi scim-m17n`
- `Japanese-yum install fonts-japanese scim-anthy`
- `Kannada-yum install fonts-kannada m17n-db-kannada scim-m17n`
- `Korean-yum install fonts-korean scim-hangul`
- `Malayalam-yum install fonts-malayalam m17n-db-malayalam scim-m17n`
- `Marathi-yum install fonts-hindi m17n-db-marathi scim-m17n`
- `Oriya-yum install fonts-oriya m17n-db-oriya scim-m17n`
- `Punjabi-yum install fonts-punjabi m17n-db-punjabi scim-m17n`
- `Sinhala-yum install fonts-sinhala m17n-db-sinhala scim-m17n`
- `Tamil-yum install fonts-tamil m17n-db-tamil scim-m17n`
- `Telugu-yum install fonts-telugu m17n-db-telugu scim-m17n`

It is also recommended that you install `scim-bridge-gtk` and `scim-qtimm` when enabling additional language support. The `scim-bridge-gtk` package prevents possible binary conflicts with third-party applications linked against older versions of `libstdc++`.

Note that additional language support packs are also available for OpenOffice (`openoffice.org-langpack-<language code>_<locale>`) and KDE (`kde-i18n-<language>`). These packages can also be installed through `yum`.

im-chooser

A new user configuration tool called `im-chooser` has been added, which allows you to easily disable or enable the usage of input methods on your desktop. So if SCIM is installed but you do not wish to run it on your desktop, you can disable it using `im-chooser`.

xinputrc

At X startup, `xinput.sh` now sources `~/.xinputrc` or `/etc/X11/xinit/xinputrc` instead of searching config files under `~/.xinput.d/` or `/etc/xinit/xinput.d/`.

Pango Support in Firefox

Firefox in Red Hat Enterprise Linux 5 is built with Pango, which provides better support for certain scripts, such as Indic and some CJK scripts.

To disable the use of Pango, set `MOZ_DISABLE_PANGO=1` in your environment before launching Firefox.

Fonts

Support is now available for synthetic emboldening of fonts that do not have a bold face.

New fonts for Chinese have been added: AR PL ShanHeiSun Uni (`uming.ttf`) and AR PL ZenKai Uni (`ukai.ttf`). The default font is AR PL ShanHeiSun Uni, which contains embedded bitmaps. If you prefer outline glyphs, add the following section to your `~/ .font.conf` file:

```
<fontconfig>
  <match target="font">
    <test name="family" compare="eq">
      <string>AR PL ShanHeiSun Uni</string>
    </test>
    <edit name="embeddedbitmap" mode="assign">
      <bool>>false</bool>
    </edit>
  </match>
</fontconfig>
```

gtk2 IM submenu

The Gtk2 context menu IM submenu no longer appears by default. You can enable it on the command line with the following command:

```
gconftool-2 --type bool --set  
'/desktop/gnome/interface/show_input_method_menu' true
```

Support for text installation on CJK

CJK (Chinese, Japanese, and Korean) rendering support has been removed from the Anaconda text installation. The text installation method is being deprecated in the long term, as the GUI installation, VNC and kickstart methods are preferred.

gtk+ deprecation

The following packages are deprecated and scheduled for removal in Red Hat Enterprise Linux:

- **gtk+**
- **gdk-pixbuf**
- **glib**

These packages are being deprecated in favor of the **gtk2** stack, which offers better functionality particularly in terms of internationalization and font handling.

CJK input on console

If you need to display Chinese, Japanese, or Korean text on the console, you need to set up a framebuffer; afterwards, install **bogl-bterm**, and run **bterm** on the framebuffer.

3.7. Kernel Notes

This section notes the differences between 2.6.9 (on which Red Hat Enterprise Linux 4 is based) and 2.6.18 (which Red Hat Enterprise Linux 5 will inherit) as of July 12, 2006. Additional features which we are currently working on upstream (for example, virtualization) that will appear late in 2.6.18 or 2.6.19 are not highlighted here. In other words, this list only shows what is already included in the upstream Linus tree; not what is currently in development. Consequently, this list is not a final, or complete list of the new Red Hat Enterprise Linux 5 features, although it does give a good overview of what can be

expected. Also, note that this section only picks out highlights of upstream changes, and as such it is not fully comprehensive. It does not include mention of several low-level hardware support enhancements and device driver info.

The following is a good source for a next level-of-detail view:

<http://kernelnewbies.org/LinuxChanges>

Performance / Scalability

- Big Kernel Lock preemption (2.6.10)
- Voluntary preemption patches (2.6.13) (subset in Red Hat Enterprise Linux 4)
- Lightweight user-space priority inheritance (PI) support for futexes, useful for real-time applications (2.6.18)
 - ref: <http://lwn.net/Articles/178253/>
- New 'mutex' locking primitive (2.6.16)
- High resolution timers (2.6.16)
 - In contrast to the low-resolution timeout API implemented in kernel/timer.c, hrtimers provide finer resolution and accuracy depending on system configuration and capabilities. These timers are currently used for itimers, POSIX timers, nanosleep and precise in-kernel timing.
- Modular, on-the-fly switchable I/O schedulers (2.6.10)
 - This was adjustable only by boot option in Red Hat Enterprise Linux 4 (also system-wide instead of per-queue).
- New Pipe implementation (2.6.11)
 - 30-90% performance improvement in pipe bandwidth
 - circular buffer allows more buffering than blocking writers
- "Big Kernel Semaphore": turns the Big Kernel Lock into a semaphore
 - reduces latency by breaking up long lock hold times and adding voluntary preemption
- X86 "SMP alternatives"
 - optimizes a single kernel image at runtime according to the available platform
 - ref: <http://lwn.net/Articles/164121/>
- kernel-headers package
 - replaces the `glibc-kernheaders` package
 - provides better suitability with the new `headers_install` feature of the 2.6.18 kernel
 - notable kernel header-related changes:
 - removed `<linux/compiler.h>` header file, as it is no longer useful

- removed `_syscallX()` macros; user-space should use `syscall()` from the C library instead
- removed `<asm/atomic.h>` and `<asm/bitops.h>` header files; C compiler provides its own atomic built-in functions better suitable for user-space programs
- content previously protected with `#ifdef __KERNEL__` is now removed completely with the `unifdef` tool; defining `__KERNEL__` in order to view parts which should not be visible to user-space is no longer effective
- removed the `PAGE_SIZE` macro from some architectures, due to variance in page sizes; user-space should be using `sysconf (_SC_PAGE_SIZE)` or `getpagesize()`
- to provide better suitability for user-space, removed several header files and header content

Generic Feature Additions

- **kexec and kdump (2.6.13)**
 - `diskdump` and `netdump` have been replaced by `kexec` and `kdump`, which ensure faster boot-up and creation of reliable kernel vmcores for diagnostic purposes. For more information and configuration instructions, please refer to `/usr/share/doc/kexec-tools-<version>/kexec-kdump-howto.txt` (replace `<version>` with the corresponding version of the `kexec-tools` package installed).
 - Note that at present, virtualized kernels cannot use the `kdump` function.
- **inotify (2.6.13)**
 - user interface for this is through the following syscalls: `sys_inotify_init`, `sys_inotify_add_watch`, and `sys_inotify_rm_watch`.
- **Process Events Connector (2.6.15)**
 - reports `fork`, `exec`, id change, and exit events for all processes to user-space.
 - Applications that may find these events useful include accounting / auditing (for example, ELSA), system activity monitoring (for example, `top`), security, and resource management (for example, CKRM). Semantics provide the building blocks for features like per-user-namespace, "files as directories" and versioned file systems.
- **Generic RTC (RealTime Clock) subsystem (2.6.17)**
- **splice (2.6.17)**
 - new IO mechanism which avoids data copies when transferring data between applications
 - ref: <http://lwn.net/Articles/178199/>

File System / LVM

- **EXT3**
 - support for Extended Attributes in the body of large inode in ext3: saves space and

improves performance in some cases (2.6.11)

- Device mapper multipath support
- ACL support for NFSv3 and NFSv4 (2.6.13)
- NFS: supports large reads and writes on the wire (2.6.16)
 - The Linux NFS client now supports transfer sizes of up to 1MB.
- VFS changes
 - The "shared subtree" patches have been merged. (2.6.15)
 - ref: <http://lwn.net/Articles/159077/>
- Big CIFS update (2.6.15)
 - features several performance improvements as well as support for Kerberos and CIFS ACL
- autofs4: updated to provide direct mount support for user-space autofs (2.6.18)
- cachefs core enablers (2.6.18)

Security

- Multilevel security implementation for SELinux (2.6.12)
- Audit subsystem
 - support for process-context based filtering (2.6.17)
 - more filter rule comparators (2.6.17)
- TCP/UDP getpeercon: enabled security-aware applications to retrieve the entire security context of a process on the other side of a socket using an IPSec security association. If only MLS-level information is needed or interoperability with legacy unix system is required, NetLabel can be used in place of IPSec.


Networking

- Added several TCP congestion modules (2.6.13)
- IPv6: supports several new sockopt / ancillary data in Advanced API (2.6.14)
- IPv4/IPv6: UFO (UDP Fragmentation Offload) Scatter-gather approach (2.6.15)
 - UFO is a feature wherein the Linux kernel network stack will offload the IP fragmentation functionality of large UDP datagram to hardware. This will reduce the overhead of stack in fragmenting the large UDP datagram to MTU-sized packets.
- Added nf_contrack subsystem (2.6.15)
 - The existing connection tracking subsystem in netfilter can only handle ipv4. There were two choices present to add connection tracking support for ipv6; either duplicate all of the ipv4 connection tracking code into an ipv6 counterpart, or (the choice taken

by these patches) design a generic layer that could handle both ipv4 and ipv6 and thus requiring only one sub-protocol (TCP, UDP, etc.) connection tracking helper module to be written. In fact, `nf_conntrack` is capable of working with any layer 3 protocol.

- IPV6
 - RFC 3484-compliant source address selection (2.6.15)
 - added support for Router Preference (RFC4191) (2.6.17)
 - added Router Reachability Probing (RFC4191) (2.6.17)
 - added support for Multiple Routing Tables and Policy Routing
- Wireless updates
 - hardware crypto and fragmentation offload support
 - QoS (WME) support, "wireless spy support"
 - mixed PTK/GTK
 - CCMP/TKIP support and WE-19 HostAP support
 - BCM43xx wireless driver
 - ZD1211 wireless driver
 - WE-20, version 20 of the Wireless Extensions (2.6.17)
 - added the hardware-independent software MAC layer, "Soft MAC" (2.6.17)
 - added LEAP authentication type
- Added generic segmentation offload (GSO) (2.6.18)
 - can improve performance in some cases, though it needs to be enabled through `ethtool`
- DCCPv6 (2.6.16)

Added Hardware Support



NOTE

This section only enumerates the most generic features among many.

- x86-64 clustered APIC support (2.6.10)
- Infiniband support (2.6.11)
- Hot plug
 - added generic memory add/remove and supporting functions for memory hotplug (2.6.15)
- SATA/libata enhancements, additional hardware support

- A completely reworked libata error handler; the result of all this work should be a more robust SATA subsystem which can recover from a wider range of errors.
- Native Command Queuing (NCQ), the SATA version of tagged command queuing - the ability to have several I/O requests to the same drive outstanding at the same time. (2.6.18)
- Hotplug support (2.6.18)
- EDAC support (2.6.16)
 - The EDAC goal is to detect and report errors that occur within the system.
- Added a new ioatdma driver for the Intel(R) I/OAT DMA engine (2.6.18)

NUMA (Non-Uniform Memory Access) / Multi-core

- Cpusets (2.6.12)
 - Cpusets now provide a mechanism for assigning a set of CPUs and Memory Nodes to a set of tasks. Cpusets constrain the CPU and memory placement of tasks only to the resources within a task's current cpuset. These are essential in managing dynamic job placement on large systems.
- NUMA-aware slab allocator (2.6.14)
 - This creates slabs on multiple nodes and manages slabs in such a way that locality of allocations is optimized. Each node has its own list of partial, free and full slabs. All object allocations for a node occur from node-specific slab lists.
- Swap migration (2.6.16)
 - Swap migration allows the moving of physical location of pages between nodes in a NUMA system while the process is running.
- Huge pages (2.6.16)
 - Added NUMA policy support for huge pages: the `huge_zonelist()` function in the memory policy layer provides a list of zones ordered by NUMA distance. The hugetlb layer will walk that list looking for a zone that has available huge pages but is also in the nodeset of the current cpuset.
 - Huge pages now obey cpusets.
- Per-zone VM counters
 - provide zone-based VM statistics, which are necessary in determining what state of memory a zone is in
- Netfilter ip_tables: NUMA-aware allocation. (2.6.16)
- Multi-core
 - Added a new scheduler domain for representing multi-core with shared caches between cores. This makes it possible to make smarter cpu scheduling decisions on such systems, improving performance greatly for some cases. (2.6.17)

- Power saving policy for the CPU scheduler: with multicore/smt cpus, the power consumption can be improved by leaving some packages idle while others do all the work, instead of spreading the tasks over all CPUs.

4. RELEASE NOTES FOR PPC

4.1. Installation-Related Notes

The following section includes information specific to the installation of Red Hat Enterprise Linux and the Anaconda installation program.



NOTE

In order to upgrade an already-installed Red Hat Enterprise Linux, you must use Red Hat Network to update those packages that have changed.

You may use Anaconda to perform a fresh installation of Red Hat Enterprise Linux 5 or to perform an upgrade from *the latest updated version* of Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5.

If you are copying the contents of the Red Hat Enterprise Linux 5 CD-ROMs (in preparation for a network-based installation, for example) be sure to copy the CD-ROMs for the operating system *only*. Do not copy the Supplementary CD-ROM, or any of the layered product CD-ROMs, as this will overwrite files necessary for Anaconda's proper operation. These CD-ROMs must be installed *after* Red Hat Enterprise Linux has been installed.

Note that the minimum RAM required to install Red Hat Enterprise Linux 5 has been raised to 1GB; the recommended RAM is 2GB. If a machine has less than 1GB RAM, the installation process may hang.

Further, PPC machines that have 1GB of RAM experience significant performance issues under certain RAM-intensive workloads. For a Red Hat Enterprise Linux 5 system to perform RAM-intensive processes optimally, it is recommended that 4GB of RAM be equipped on the machine. This ensures that the system has the same number of physical pages as that on PPC machines (using 512MB of RAM) with previous versions of Red Hat Enterprise Linux installed.

4.1.1. ISO Contents and Registration

The organization of software component packages into product specific-variants has changed from previous versions of Red Hat Enterprise Linux. The number of variants and ISO images has been reduced to only one: the Red Hat Enterprise Linux 5 Server.

In the future, additional repositories for a number of options that provide added functionality over the core Server distribution may be provided.

With optional content in the same tree or ISO image, it is important to avoid a mismatch between the components offered for installation and those covered by the subscription. Such a mismatch could result in an increased exposure to bug and vulnerability risks.

In order to ensure that the components offered are in sync with the subscription, Red Hat Enterprise Linux 5 requires entering an *Installation Number* that will be used to configure the installer to offer the right package set. This Installation Number is included in your subscription.

If you skip entering the Installation Number, this will result in a core Server or Desktop installation. Additional functionality can then be added manually at a later time. For more information about Installation Numbers, please refer to <http://www.redhat.com/apps/support/in.html>.

The Installation Number used during the installation process will be saved in `/etc/sysconfig/rhn/install-num`. When registering with Red Hat Network, this file will be referenced by `rhn_register` to automatically determine which appropriate child channels the system should be subscribed to.

4.1.2. New RPM GPG Signing Keys

A new release signing key is used to sign Red Hat Enterprise Linux 5 packages. When updating a system for the first time, you will be prompted to allow this key to be installed.

Signing keys are distributed in the following files:

- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release` – contains the public key for the new release signing key
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-auxiliary` – contains the public key for an auxiliary release signing key, currently not in use
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-former` – contains the public key for the previous release signing key, used for past Red Hat Enterprise Linux releases

4.1.3. Subversion

In Red Hat Enterprise Linux 5, the Subversion version control system is linked against Berkeley DB 4.3. If you are upgrading from Red Hat Enterprise Linux 4 and any Subversion repositories which use the Berkeley DB backend "BDB" (rather than the pure file system-based "FSFS" backend) have been created on the system, special care must be taken to ensure the repositories can be accessible after the upgrade. The following process must be performed on the Red Hat Enterprise Linux 4 system, *prior* to upgrading to Red Hat Enterprise Linux 5:

1. Shut down any running processes and ensure that no processes can access the repository (for example, `httpd`, `svnserve` or any local users with direct access).
2. Create a backup of the repository using the following command:

```
svnadmin dump /path/to/repository | gzip > repository-backup.gz
```

3. Run the `svnadmin recover` command on the repository:

```
svnadmin recover /path/to/repository
```

4. Delete any unused log files in the repository:

```
svnadmin list-unused-dblogs /path/to/repository | xargs rm -vf
```

5. Delete any remaining shared-memory files in the repository:

```
rm -f /path/to/repository/db/___db.0*
```

4.1.4. Other Installation Notes

- If IDE/PATA (*Parallel ATA*) devices are configured in "100% Native" mode, some BIOSes may prevent the Red Hat Enterprise Linux 5 installation process from completing successfully. To prevent this from occurring, configure the IDE/PATA mode as "Legacy" in the BIOS.
- The *IBM System z* does not provide a traditional Unix-style physical console. As such, Red Hat Enterprise Linux 5 for the *IBM System z* does not support the *firstboot* functionality during initial program load.

To properly initialize setup for Red Hat Enterprise Linux 5 on the *IBM System z*, run the following commands after installation:

- `/usr/bin/setup` – provided by the `setuptools` package
- `/usr/bin/rhn_register` – provided by the `rhn-setup` package
- When booting Anaconda with PXE using the parameter `ksdevice=bootif`, you will still be prompted for the ethernet interface to use during installation. If only one ethernet device is plugged in, use the parameter `ksdevice=link` instead. Alternatively, you can also specify the interface manually.

4.2. Technology Previews

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 5 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. Erratas 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.

Stateless Linux

Included in this release of Red Hat Enterprise Linux 5 are enabling infrastructure pieces for *Stateless Linux*. Stateless Linux is a new way of thinking about how a system is to be run and managed, designed to simplify provisioning and management of large numbers of systems by making them easily replaceable. This is accomplished primarily by establishing prepared system images which get replicated and managed across a large number of stateless systems, running the operating system in a read-only manner (please refer to `/etc/sysconfig/readonly-root` for more details).

In its current state of development, the Stateless features are subsets of the intended goals. As such, the capability is being labeled as a technology preview.

The following is a list of the initial capabilities included in Red Hat Enterprise Linux 5:

- running a stateless image over NFS
- running a stateless image via loopback over NFS
- running on iSCSI

It is highly recommended that those interested in testing stateless code read the HOWTO at <http://fedoraproject.org/wiki/StatelessLinuxHOWTO> and join stateless-list@redhat.com.

GFS2

GFS2 is an evolutionary advancement based on the GFS file system. While fully functional, GFS2 is not yet considered production-ready. GFS2 is targeted to become fully supported in a subsequent Red Hat Enterprise Linux 5 update. There is also an in-place conversion utility, `gfs2_convert`, which can update the metadata of the older GFS file system format, converting it to a GFS2 file system.

FS-Cache

FS-Cache is a local caching facility for remote file systems that allows users to cache NFS data on a locally mounted disk. To set up the FS-Cache facility, install the `cachefilesd` RPM and refer to the instructions in `/usr/share/doc/cachefilesd-<version>/README`.

Replace `<version>` with the corresponding version of the `cachefilesd` package installed.

Compiz

Compiz is an OpenGL-based compositing window manager. In addition to regular window management, `compiz` also acts as a compositing manager, coordinating and synchronizing the overall desktop redrawing to provide a smoother desktop experience with less flicker.

`Compiz` uses 3D hardware acceleration to render effects such as live thumbnail windows, window drop shadows, animated window minimizing and transitions between virtual desktops.

Due to limitations in the current rendering architecture, `compiz` cannot work correctly with direct rendering OpenGL applications or applications using the Xv extension. Such applications will exhibit harmless rendering artifacts; because of this, `compiz` is currently a technology preview.

Enhancement for Ext3

In Red Hat Enterprise Linux 5, the EXT3 file system capacity has been extended beyond 8TB to a maximum of 16TB. This capability is being included as a technology preview, and is targeted for full support in a future release of Red Hat Enterprise Linux 5.

AIGLX

AIGLX is a technology preview feature of the otherwise fully supported X server. It aims to enable GL-accelerated effects on a standard desktop. The project consists of the following:

- a lightly modified X server
- an updated Mesa package that adds new protocol support

By installing these components, you can have GL-accelerated effects on your desktop with very few changes, as well as the ability to enable and disable them at will without replacing your X server. AIGLX also enables remote GLX applications to take advantage of hardware GLX acceleration.

Frysk GUI

The goal of the *frysk* project is to create an intelligent, distributed, always-on system monitoring and debugging tool that allows developers and system administrators to:

- monitor running processes and threads (including creation and destruction events)
- monitor the use of locking primitives

- expose deadlocks
- gather data
- debug any given process by choosing it from a list or allowing `frysk` to open a source code (or other) window on a process that is crashing or misbehaving

In Red Hat Enterprise Linux 5 the `frysk` graphical user interface is a technology preview, whereas the `frysk` command line interface is fully supported.

Systemtap

Systemtap provides free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. With the help of `systemtap`, developers no longer need to go through the tedious and disruptive instrument, recompile, install, and reboot sequence that may be otherwise required to collect data.

Dogtail

Dogtail is a GUI test tool and automation framework written in Python that uses *Accessibility* technologies to communicate with desktop applications.

Support for Indic Languages and Sinhalese

Red Hat Enterprise Linux 5 also features support for the following languages as technology preview:

- *Assamese*
- *Kannada*
- *Sinhalese*
- *Telugu*

For more information about how to install and enable support for these languages, refer to the *Internationalization* section of this document.

Installing to dm-multipath Devices

Anaconda now has the capability to detect, create, and install to *dm-multipath* devices. To enable this feature, add the parameter `mpath` to the kernel boot line.

Note that the parameter `mpath` may cause a boot failure if a device's `major:minor` number changes. This issue will be addressed in a future update of Red Hat Enterprise Linux 5.

Installation / Boot for iSCSI software initiator (open-iscsi)

Anaconda now provides the ability to install to an iSCSI device. Booting and installing is fully supported with the *QLogic qla4xxx* hardware initiator. However, the capability to install to an iSCSI device for the `open-iscsi` software initiator is currently considered a Technology Preview, due to the following issues:

- Text mode installation does not complete. You must do a graphical install, or an automated kickstart install.
- Media-based installations do not complete. You must do a network-based install.
- Depending on the timing of events, Anaconda may be unable to detect all the iSCSI targets

or LUNs. When this occurs, use the installer shell to configure the storage through iSCSI commands.

- The `iscsid` daemon may not properly start. Such an occurrence will prevent the system from handling all iSCSI errors, such as network problems, SCSI/iSCSI timeouts, and target errors. To confirm that the `iscsid` daemon is running, run the command `iscsiadm -m session -i` and check that the line, **Internal iscsid Session State:** reports a value (it can be any value).
- On certain iSCSI target implementations, the system may hang during shutdown.
- On certain iSCSI target implementations, the system may hang during reboot. To avoid this, shutdown the system and boot it up again (instead of rebooting directly from a session).
- Booting from iSCSI devices on the *IBM System p* does not work reliably. While installation on an iSCSI device may appear to succeed, the resulting installation will not boot properly.
- On the first boot after install, you may receive SELinux errors such as the following:

```
kernel: audit(1169664832.270:4): avc: denied { read
} for pid=1964 comm="iscsid"
```

To work around this, boot the system with the kernel parameter `enforcing=0`. Once the system has properly booted, use the command `setenforce 1` to restore enforcing mode.

These limitations will be addressed in a future Red Hat Enterprise Linux 5 update.

4.3. Known Issues

- Host bus adapters that use the MegaRAID driver must be set to operate in "Mass Storage" emulation mode, not in "I2O" emulation mode. To do this, perform the following steps:
 1. Enter the MegaRAID BIOS **Set Up Utility**.
 2. Enter the **Adapter settings** menu.
 3. Under **Other Adapter Options**, select **Emulation** and set it to **Mass Storage**.

If the adapter is incorrectly set to "I2O" emulation, the system will attempt to load the i2o driver. This will fail, and prevent the proper driver from being loaded.

Previous Red Hat Enterprise Linux releases generally do not attempt to load the I2O driver before the MegaRAID driver. Regardless of this, the hardware should always be set to "Mass Storage" emulation mode when used with Linux.

- Red Hat Enterprise Linux 5 includes `openmpi-1.1.1-4.e15` (from the OFED 1.1 distribution), which has been discovered to eventually quit working entirely. This happens after the `openmpi` stack works as expected for a varying amount of time.

For updated versions of `openmpi`, please check <http://people.redhat.com/dledford/Infiniband/openmpi>

- Installing Windows Server 2003 as a guest on a fully virtualized Red Hat Enterprise Linux 5 system ends unexpectedly after completing the first stage of installation. When this happens, the graphical console window closes, and the guest disappears from the **Virtual Machine**

Manager's list of machines, resulting in a **Broken pipe** error.

This issue will be resolved in an upcoming Red Hat Enterprise Linux 5 update. To work around this, use the following command at the terminal:

```
xm create /etc/xen/<name of guest machine>
```

Afterwards, open the virtual machine.

- When attempting to create a fully virtualized Windows Server 2003 from a CD / DVD, the second stage of the guest install will not continue upon reboot.

To work around this, edit `/etc/xen/<name of guest machine>` by properly appending an entry for the CD / DVD device.

If an installation to a simple file is used as a virtual device, the `disk` line of `/etc/xen/<name of guest machine>` will read like the following:

```
disk = [ 'file:/PATH-OF-SIMPLE-FILE,hda,w' ]
```

A DVD-ROM device located on the host as `/dev/dvd` can be made available to stage 2 of the installation as `hdc` by appending an entry like `'phy:/dev/dvd,hdc:cdrom,r'`. As such, the disk line should now read as follows:

```
disk = [ 'file:/opt/win2003-sp1-20061107,hda,w',  
        'phy:/dev/dvd,hdc:cdrom,r' ]
```

The precise device path to use may vary depending on your hardware.

- `rmmod xennet` causes `domU` to crash; this is caused by a grant table issue in the Virtualization feature. Due to the current inability of the Virtualization feature to asynchronously release grant table operations, it is unsafe to unload the `xennet` module in guests. In such situations, grant tables are used to perform backend-frontend communication, and there is no guarantee that the backend will release the references, leading to an inevitable memory leak.

This issue will be resolved in the next minor release of Red Hat Enterprise Linux 5. At present, users are advised *not* to unload the `xennet` module in guests.

- Running `ethtool eth0` outputs incomplete information about the ethernet card settings. This only occurs in systems running a virtualized kernel, since the Virtualization feature uses a networking setup where the physical ethernet device is identified as `peth0`. As such, the correct command for retrieving information about the physical ethernet device is `ethtool peth0`.
- When using Red Hat Enterprise Linux 5 on a machine with an *nVidia CK804* chipset installed, you may receive kernel messages similar to the following:

```
kernel: assign_interrupt_mode Found MSI capability  
kernel: pcie_portdrv_probe->Dev[005d:10de] has invalid IRQ. Check  
vendor BIOS
```

These messages indicate that certain PCI-E ports are not requesting IRQs. Further, these messages do not, in any way, affect the operation of the machine.

- Some *Cisco Aironet Wireless* devices prevent **NetworkManager** from storing connection details for wireless networks that do not broadcast an SSID. This is caused by a *Cisco Aironet Wireless* device firmware limitation.
- Laptops that have the *Cisco Aironet MPI-350* wireless card equipped may hang trying to get a DHCP address during any network-based installation using the wired ethernet port.

To work around this, use local media for your installation. Alternatively, you can disable the wireless card in the laptop BIOS prior to installation (you can re-enable the wireless card after completing the installation).

- Currently, `system-config-kickstart` does not support package selection and deselection. When using `system-config-kickstart`, the **Package Selection** option indicates that it is disabled. This is because `system-config-kickstart` uses `yum` to gather group information, but is unable to configure `yum` to connect to Red Hat Network.

This issue is currently being investigated for resolution by the next minor release of Red Hat Enterprise Linux 5. At present, you need to update package sections in your kickstart files manually. When using `system-config-kickstart` to open a kickstart file, it will preserve all package information in it and write it back out when you save.

- 4-socket AMD *Sun Blade X8400 Server Modules* systems that do not have memory configured in `node 0` will panic during boot. Systems should be configured with memory in `node 0` to prevent the kernel panic.
- When using `Alt-SysRq-W` to debug, the following warning message will appear:

```
Badness in smp_call_function at arch/powerpc/kernel/smp.c:223
```

Afterwards, the system will also warn that it will hang. This message should be ignored as it will not cause the system to hang.

- Installing to LVM mirror devices through Anaconda is currently not supported. This capability will be added in a future update of Red Hat Enterprise Linux 5.
- When installing Red Hat Enterprise Linux 5 from a directory on an NFS server containing Red Hat Enterprise Linux ISO images, Anaconda may display the following error message:

```
Unable to read package metadata. This may be due to a missing
repopdata directory.
Please ensure that your install tree has been correctly generated.
Cannot
open/read repomd.xml file for repository:
```

This problem occurs if the directory holding the ISO images also contains a partially unpacked installation tree (for example, the `/images` directory from the first ISO). The presence of such directories results in the error stated above.

To prevent this error, unpack trees only to directories other than the one containing the installation ISO images.

- Boot-time logging to `/var/log/boot.log` is not available in this release of Red Hat Enterprise Linux 5. An equivalent functionality will be added in a future update of Red Hat Enterprise Linux 5.
- Neither `kexec` nor `kdump` are able to dump onto disks attached to an *accraid* controller.

To work around this issue, use `scp` for network dumping. Alternatively, you can also dump onto a disk through a different controller.

- Installing a fully virtualized guest using split installation media -- specifically, multiple CD-ROMs -- may fail when required to switch between installation CDs. During the guest OS installation process, users may be prevented from mounting or ejecting installation CDs, which prevents the installation from completing.

As such, it is recommended that you use the QEMU monitor console to switch CD-ROM images during the guest OS installation process. The procedure is as follows:

1. Open a graphical VNC console to the guest OS.
2. Unmount the CD-ROM device in the guest OS.
3. Switch to the QEMU monitor console by pressing `Ctrl-Alt-2`.
4. Run the command `eject hdc`.
5. Run the command `change hdc <path to the CD-ROM in host system>`.
6. Switch back to the guest OS console by pressing `Ctrl-Alt-1`.
7. Mount the CD-ROM device in the guest OS.

Note that when using a regular VNC client the host X server may encounter some difficulty interpreting the `Ctrl-Alt-2` and `Ctrl-Alt-1` command. To work around this in `virt-manager`, use *sticky keys*. Pressing `Ctrl` three times makes it "sticky" until the next non-modifier is pressed. As such, to send `Ctrl-Alt-1`, press `Ctrl` twice before pressing `Ctrl-Alt-1`.

- The Red Hat Enterprise Linux 5 *Driver Update Model* creates modified `initrd` images whenever a `kmod` package that includes a bootpath-modifying driver is installed. In time, the number of backup `initrd` images may soon fill the `/boot` partition, particularly if the system undergoes a sizeable number of driver updates.

As such, it is recommended that you monitor the free space on the `/boot` partition if you regularly perform driver updates. You can free up more space in `/boot` by removing older `initrd` images; these files end in `.img0`, `.img1`, `.img2`, and so on.

- Autorun on removable media is currently disabled. To install packages from the Red Hat Enterprise Linux Supplementary CD, launch the CD installer manually using the following command:

```
system-cdinstall-helper /media/path-to-mounted-drive
```

- When upgrading from Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5, the Deployment Guide is not automatically installed. You need to use `pirut` to manually install it after completing the upgrade.
- An `autofs` bug prevents multi-mounts from working properly.

During an expiry, if the last multi-mount component to be checked does not have a mount associated with it while other components are busy, `autofs` erroneously determines the multi-mount to be expirable. This causes the multi-mount to be partially expired, resulting in

the multi-mount becoming unresponsive to further mount requests and expire runs.

To permanently resolve this problem, update `autofs` using the command `yum update autofs`.

- The system may not successfully reboot into a `kexec/kdump` kernel if X is running and using a driver other than `vesa`. This problem only exists with *ATI Rage XL* graphics chipsets.

If X is running on a system equipped with *ATI Rage XL*, ensure that it is using the `vesa` driver in order to successfully reboot into a `kexec/kdump` kernel.

4.4. General Information

This section contains general information not specific to any other section of this document.

Red Hat Enterprise Linux Deployment Guide

This release of Red Hat Enterprise Linux includes a fairly comprehensive Deployment Guide. To access it, go to **System** (on the top panel) => **Documentation** => **Red Hat Enterprise Linux Deployment Guide**.

It is the intention of Red Hat to provide fully localized versions of the Deployment Guide for all supported languages. If you have installed a localized version of the Deployment Guide, it is recommended that you update it when a new version becomes available through Red Hat Network.

Web Server Packaging Changes

Red Hat Enterprise Linux 5 now includes version 2.2 of the Apache HTTP Server. This release brings a number of improvements over the 2.0 series, including:

- improved caching modules (`mod_cache`, `mod_disk_cache`, `mod_mem_cache`)
- a new structure for authentication and authorization support, replacing the authentication modules provided in previous versions
- support for proxy load balancing (`mod_proxy_balancer`)
- support for handling large files (namely, greater than 2GB) on 32-bit platforms

The following changes have been made to the default `httpd` configuration:

- The `mod_cern_meta` and `mod_asis` modules are no longer loaded by default.
- The `mod_ext_filter` module is now loaded by default.

If you are upgrading from a previous release of Red Hat Enterprise Linux, the `httpd` configuration will need to be updated for `httpd 2.2`. For more information, refer to <http://httpd.apache.org/docs/2.2/upgrading.html>.

Note that any third-party modules compiled for `httpd 2.0` must be rebuilt for `httpd 2.2`.

php

Version 5.1 of PHP is now included in Red Hat Enterprise Linux 5, which includes a number of changes to the language along with significant performance improvements. Some scripts might need to be edited for use with the new version; please refer to the link below for more information on migrating from PHP 4.3 to PHP 5.1:

<http://www.php.net/manual/en/migration5.php>

The `/usr/bin/php` executable is now built using the CLI command-line SAPI, rather than the CGI SAPI. Use `/usr/bin/php-cgi` for CGI SAPI. The `php-cgi` executable also includes FastCGI support.

The following extension modules have been added:

- the `mysqli` extension, a new interface designed specifically for MySQL 4.1 (included in the `php-mysql` package)
- `date`, `hash`, `Reflection`, `SPL` and `SimpleXML` (built-in with the `php` package)
- `pdo` and `pdo_sqlite` (in the `php-pdo` package)
- `pdo_mysql` (in the `php-mysql` package)
- `pdo_pgsql` (in the `php-pgsql` package)
- `pdo_odbc` (in the `php-odbc` package)
- `soap` (in the `php-soap` package)
- `xmlreader` and `xmlwriter` (in the `php-xml` package)
- `dom` (replacing the `domxml` extension in the `php-xml` package)

The following extension modules are no longer included:

- `dbx`
- `dio`
- `yp`
- `overload`
- `domxml`

The PEAR Framework

The PEAR framework is now packaged in the `php-pear` package. Only the following PEAR components are included in Red Hat Enterprise Linux 5:

- `Archive_Tar`
- `Console_Getopt`
- `XML_RPC`

Encrypted Swap Partitions and Non-root File Systems

Red Hat Enterprise Linux 5 now provides basic support for encrypted swap partitions and non-root file systems. To use these features, add the appropriate entries to `/etc/crypttab` and reference the created devices in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry:

–

```
my_swap /dev/hdb1 /dev/urandom swap,cipher=aes-cbc-essiv:sha256
```

This creates the encrypted block device `/dev/mapper/my_swap`, which can be referenced in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry for a file system volume:

```
my_volume /dev/hda5 /etc/volume_key cipher=aes-cbc-essiv:sha256
```

The `/etc/volume_key` file contains a plaintext encryption key. You can also specify `none` as the key file name; this configures the system to ask for the encryption key during boot instead.

It is recommended to use LUKS (Linux Unified Key Setup) for setting up file system volumes. To do this, follow these steps:

1. Create the encrypted volume using `cryptsetup luksFormat`.
2. Add the necessary entry to `/etc/crypttab`.
3. Set up the volume manually using `cryptsetup luksOpen` (or reboot).
4. Create a file system on the encrypted volume.
5. Add the necessary entry to `/etc/fstab`.

mount and umount

The `mount` and `umount` commands no longer directly support NFS; a built-in NFS client no longer exists. A separate `nfs-utils` package, which provides `/sbin/mount.nfs` and `/sbin/umount.nfs` helpers, must be installed for this.

CUPS Printer Browsing

CUPS printer browsing over a local subnet can be configured using the graphical tool `system-config-printer`. It can also be done using the CUPS web interface, <http://localhost:631/>.

To use directed broadcasts for printer browsing between subnets, open `/etc/cups/cupsd.conf` on the clients and replace `BrowseAllow @LOCAL` with `BrowseAllow ALL`.

ATI and R500 Support

ATI graphics cards based on the R500 chipset are supported for the `vesa` driver only, and are not supported by Red Hat Enterprise Linux 5 on external monitors, LCD projectors or accelerated 3D support.

up2date and yum

`up2date` is being deprecated in favor of `yum` (Yellowdog Updater Modified). As such, it is advisable that you revise any `up2date`-dependent scripts your system is using accordingly. For more information about `yum`, consult its man page with the command `man yum`; you can also consult the installed documentation under the directories `/usr/share/doc/yum-<version>` and `/usr/share/doc/yum-metadata-parser-<version>` (replace `<version>` with the corresponding version of `yum` and `yum-metadata-parser` installed).

OpenLDAP Server and Red Hat Directory Server

Red Hat Directory Server is an LDAP-based server that centralizes enterprise and network data into an OS-independent, network-based registry. It is set to replace OpenLDAP server components, which will be deprecated *after* Red Hat Enterprise Linux 5. For more information about Red Hat Directory Server, refer to <http://www.redhat.com/software/rha/directory/>.

i810 Driver and i830 Support

The *i810* driver supports all integrated Intel graphics chipsets, from i810 to i965. However, the support for i830 (and newer) chipsets is limited; the i810 driver can only set modes listed in the video BIOS. If your machine has an i830 or newer chipset installed, run the following command to determine what the available modes are:

```
grep Mode: /var/log/Xorg.0.log
```

Modes marked with an asterisk (*) are available for selection.

Many laptop video BIOSes do not supply a mode that matches the native panel size. Therefore the chosen mode may appear stretched, distorted, or with black borders. As such, if your chosen mode does not display properly, you need a BIOS update from your hardware vendor for the native panel size to work correctly.

Intel PRO/Wireless 3945ABG Network Connection Support

This release of Red Hat Enterprise Linux 5 includes support for the *ipw3945* (*Intel PRO/Wireless 3945ABG Network Connection*) adapter. The Red Hat Enterprise Linux 5 Supplementary disc contains the driver, regulatory daemon and firmware needed to support this adapter.

To enable support for the ipw3945 wireless adapter, search the Red Hat Enterprise Linux 5 Supplementary disc for packages with filenames containing "3945" and install them.

rawio

rawio is a deprecated interface; however, Red Hat Enterprise Linux 5 still includes support for it. If you have an application that performs device access using *rawio*, it is highly recommended that you modify your application to open the block device with the `O_DIRECT` flag. The *rawio* interface will remain throughout the life of Red Hat Enterprise Linux 5, but is a candidate for removal in a future release.

Currently, AIO (*Asynchronous I/O*) on file systems is only supported in `O_DIRECT` or non-buffered mode. Further, note that the asynchronous poll interface is no longer present, and that AIO on pipes is no longer supported.

ctmpc

ctmpc is a deprecated driver; however, it will still be included throughout the life of Red Hat Enterprise Linux 5. Note that it is a candidate for removal from future releases.

Policy Modules and semanage Support

Red Hat Enterprise Linux 5 now supports *policy modules* and *semanage*. Policy modules simplify the creation and distribution of policy customizations and third-party policies through the use of the `semodule` and `checkmodule` tools.

The *semanage* tool is a policy management tool that modifies the SELinux configuration. It also allows you to configure file contexts, networking component labeling, and user mappings for Linux-to-SELinux.

raw Device Mapping

The raw devices interface has been deprecated in Red Hat Enterprise Linux 5; raw device mapping is now configured via `udev` rules.

To configure raw device mapping, add the appropriate entries to `/etc/udev/rules.d/60-raw.rules` in the following formats:

- For device names:

```
ACTION=="add", KERNEL="device name", RUN+="raw /dev/raw/rawX %N"
```

- For major / minor numbers:

```
ACTION=="add", ENV{MAJOR}="A", ENV{MINOR}="B", RUN+="raw
/dev/raw/rawX %M %m"
```

Replace `<device name>` with the name of the device you need to bind (for example, `/dev/sda1`). "A" and "B" are the major / minor numbers of the device you need to bind, and `X` is the raw device number that you want the system to use.

If you have a large, pre-existing `/etc/sysconfig/rawdevices` file, convert it with the following script:

```
#!/bin/sh

grep -v "^ *#" /etc/sysconfig/rawdevices | grep -v "^$" | while read dev
major minor ; do
  if [ -z "$minor" ]; then
    echo "ACTION==\"add\", KERNEL==\"${major##/dev/}\",
RUN+=\"/usr/bin/raw $dev %N\""
  else
    echo "ACTION==\"add\", ENV{MAJOR}==\"$major\", ENV{MINOR}==\"$minor\",
RUN+=\"/usr/bin/raw $dev %M %m\""
  fi
done
```

QLogic Support

Red Hat Enterprise Linux 5 supports the QLogic family of iSCSI HBA (*Host Bus Adapters*). At present, only the iSCSI interface to these boards are supported (using the `qla4xxx` driver).

In addition, Red Hat does not currently support these boards as Ethernet NIC, as this capability requires the `qla3xxx` driver. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.

IBM System z Instruction Set

In order to optimally exploit the *IBM System z* instruction set for 31-bit applications, it is recommended that you use the `gcc` option `-march=z900`. For 64-bit applications, the `gcc` will exploit the *IBM System z* instruction set by default.

iSeries Access for Linux

The *iSeries ODBC Driver* for Linux has been replaced by the *iSeries Access for Linux*, which can be downloaded at the following link:

<http://www.ibm.com/eserver/series/access/linux/>

The *iSeries Access for Linux* offers Linux-based access to iSeries servers, and allows you to:

- Access the DB2 UDB (*Universal Database*) for iSeries using its ODBC Driver
- Establish a 5250 session to an iSeries server from a Linux client
- Access the DB2 UDB via the EDRS (*Extended Dynamic Remote SQL*) driver
- Support 32-bit (i386 and PowerPC) and 64-bit (x86-64 and PowerPC) platforms

IBM Power4 iSeries

Red Hat Enterprise Linux no longer supports the *IBM Power4 iSeries*.

4.5. Driver Update Program

This section includes information on the implementation of the Red Hat Enterprise Linux 5 Driver Update Program.

Kernel Module Packages

On Red Hat Enterprise Linux 5, it is possible to build updated kernel module packages that depend upon the current kernel ABI version and not on a specific kernel release number. This facilitates building kernel modules that can be used against a range of Red Hat Enterprise Linux 5 kernels, rather than a single release. The project website at <http://www.kerneldrivers.org/> contains more information about the packaging process, as well as several examples.

Note that the following issues have also been identified:

- A minor ABI compatibility issue affecting the network `rhe15_drivers_s390_net_ga` symbol grouping on s390x kernels has been identified. Symbols beginning with `iucv_` should not be used by third party drivers as they may be updated sometime throughout the life of Red Hat Enterprise Linux 5. Similarly, on ppc64, the structure `spu_priv1_ops` should not be used by third parties.
- Bootpath drivers distributed as `kmod` packages are not officially supported.
- Overriding existing in-kernel drivers are not currently supported.

These issues will be addressed in a future update of Red Hat Enterprise Linux 5.

Kernel Module Loading

The module loading behavior on Red Hat Enterprise Linux 5 has changed from previous releases of Red Hat Enterprise Linux. The modules shipped in the Red Hat Enterprise Linux 5 kernel package are signed, as was the case in Red Hat Enterprise Linux 4. On Red Hat Enterprise Linux 5 kernels, however, it is no longer possible to load a signed module from another kernel build.

This means that a module shipped with the initial Red Hat Enterprise Linux 5 distribution cannot be loaded in future updated kernels. This helps prevent users from loading unsupported modules on a system. Red Hat only supports modules that are signed and included in a distribution.

If you want to load an older module, you can try rebuilding it without a signature. Alternatively, you can remove the signature from the binary file using the following command:

```
objcopy -R .module_sig <module name>-mod.ko <module name>-nosig.ko
```

It is recommended that you consult with a designated Red Hat Global Support Representative before attempting to load unsigned modules.

4.6. Internationalization

This section includes information on language support under Red Hat Enterprise Linux 5.

Input Methods

SCIM (Smart Common Input Method) has replaced IIMF as the input method system for Asian and other languages in this release. The default GTK Input Method Module for SCIM is provided by `scim-bridge`; in Qt, it is provided by `scim-qtim`.

Below are the default trigger hotkeys for different languages:

- All languages: `Ctrl-Space`
- Japanese: `Zenkaku-Hankaku` or `Alt-``
- Korean: `Shift-Space`

If SCIM is installed, it runs by default for all users.

After installing or removing SCIM engine packages, it is recommended to start a new desktop session in order for the changes to be reflected in the SCIM language menu.

Language Installation

To enable additional language support for some Asian languages, you need to install the necessary language support packages. Below is a list of these languages and the command you need to run (as root) to install their corresponding language support packages:

- Assamese – `yum install fonts-bengali m17n-db-assamese scim-m17n`
- Bengali – `yum install fonts-bengali m17n-db-bengali scim-m17n`
- Chinese – `yum install fonts-chinese scim-chewing scim-pinyin scim-tables-chinese`
- Gujarati – `yum install fonts-gujarati m17n-db-gujarati scim-m17n`
- Hindi – `yum install fonts-hindi m17n-db-hindi scim-m17n`
- Japanese – `yum install fonts-japanese scim-anthy`
- Kannada – `yum install fonts-kannada m17n-db-kannada scim-m17n`
- Korean – `yum install fonts-korean scim-hangul`
- Malayalam – `yum install fonts-malayalam m17n-db-malayalam scim-m17n`
- Marathi – `yum install fonts-hindi m17n-db-marathi scim-m17n`
- Oriya – `yum install fonts-oriya m17n-db-oriya scim-m17n`

- `Punjabi-yum install fonts-punjabi m17n-db-punjabi scim-m17n`
- `Sinhala-yum install fonts-sinhala m17n-db-sinhala scim-m17n`
- `Tamil-yum install fonts-tamil m17n-db-tamil scim-m17n`
- `Telugu-yum install fonts-telugu m17n-db-telugu scim-m17n`

It is also recommended that you install `scim-bridge-gtk` and `scim-qtimm` when enabling additional language support. The `scim-bridge-gtk` package prevents possible binary conflicts with third-party applications linked against older versions of `libstdc++`.

Note that additional language support packs are also available for OpenOffice (`openoffice.org-langpack-<language code>_<locale>`) and KDE (`kde-i18n-<language>`). These packages can also be installed through `yum`.

im-chooser

A new user configuration tool called `im-chooser` has been added, which allows you to easily disable or enable the usage of input methods on your desktop. So if SCIM is installed but you do not wish to run it on your desktop, you can disable it using `im-chooser`.

xinputrc

At X startup, `xinput.sh` now sources `~/.xinputrc` or `/etc/X11/xinit/xinputrc` instead of searching config files under `~/.xinput.d/` or `/etc/xinit/xinput.d/`.

Pango Support in Firefox

Firefox in Red Hat Enterprise Linux 5 is built with Pango, which provides better support for certain scripts, such as Indic and some CJK scripts.

To disable the use of Pango, set `MOZ_DISABLE_PANGO=1` in your environment before launching Firefox.

Fonts

Support is now available for synthetic emboldening of fonts that do not have a bold face.

New fonts for Chinese have been added: AR PL ShanHeiSun Uni (`uming.ttf`) and AR PL ZenKai Uni (`ukai.ttf`). The default font is AR PL ShanHeiSun Uni, which contains embedded bitmaps. If you prefer outline glyphs, add the following section to your `~/.font.conf` file:

```
<fontconfig>
  <match target="font">
    <test name="family" compare="eq">
      <string>AR PL ShanHeiSun Uni</string>
    </test>
    <edit name="embeddedbitmap" mode="assign">
      <bool>>false</bool>
    </edit>
  </match>
</fontconfig>
```

gtk2 IM submenu

The Gtk2 context menu IM submenu no longer appears by default. You can enable it on the command line with the following command:

```
gconftool-2 --type bool --set
'/desktop/gnome/interface/show_input_method_menu' true
```

Support for text installation on CJK

CJK (Chinese, Japanese, and Korean) rendering support has been removed from the Anaconda text installation. The text installation method is being deprecated in the long term, as the GUI installation, VNC and kickstart methods are preferred.

gtk+ deprecation

The following packages are deprecated and scheduled for removal in Red Hat Enterprise Linux:

- **gtk+**
- **gdk-pixbuf**
- **glib**

These packages are being deprecated in favor of the **gtk2** stack, which offers better functionality particularly in terms of internationalization and font handling.

CJK input on console

If you need to display Chinese, Japanese, or Korean text on the console, you need to set up a framebuffer; afterwards, install **bogl-bterm**, and run **bterm** on the framebuffer.

4.7. Kernel Notes

This section notes the differences between 2.6.9 (on which Red Hat Enterprise Linux 4 is based) and 2.6.18 (which Red Hat Enterprise Linux 5 will inherit) as of July 12, 2006. Additional features which we are currently working on upstream (for example, virtualization) that will appear late in 2.6.18 or 2.6.19 are not highlighted here. In other words, this list only shows what is already included in the upstream Linus tree; not what is currently in development. Consequently, this list is not a final, or complete list of the new Red Hat Enterprise Linux 5 features, although it does give a good overview of what can be expected. Also, note that this section only picks out highlights of upstream changes, and as such it is not fully comprehensive. It does not include mention of several low-level hardware support enhancements and device driver info.

The following is a good source for a next level-of-detail view:

<http://kernelnewbies.org/LinuxChanges>

Performance / Scalability

- Big Kernel Lock preemption (2.6.10)
- Voluntary preemption patches (2.6.13) (subset in Red Hat Enterprise Linux 4)
- Lightweight user-space priority inheritance (PI) support for futexes, useful for real-time applications (2.6.18)
 - ref: <http://lwn.net/Articles/178253/>

- New 'mutex' locking primitive (2.6.16)
- High resolution timers (2.6.16)
 - In contrast to the low-resolution timeout API implemented in kernel/timer.c, hrtimers provide finer resolution and accuracy depending on system configuration and capabilities. These timers are currently used for itimers, POSIX timers, nanosleep and precise in-kernel timing.
- Modular, on-the-fly switchable I/O schedulers (2.6.10)
 - This was adjustable only by boot option in Red Hat Enterprise Linux 4 (also system-wide instead of per-queue).
- Conversion to 4-level page tables (2.6.11)
 - allows x86-64 to increase from 512G to 128TB of memory
- New Pipe implementation (2.6.11)
 - 30-90% performance improvement in pipe bandwidth
 - circular buffer allows more buffering than blocking writers
- "Big Kernel Semaphore": turns the Big Kernel Lock into a semaphore
 - reduces latency by breaking up long lock hold times and adding voluntary preemption
- X86 "SMP alternatives"
 - optimizes a single kernel image at runtime according to the available platform
 - ref: <http://lwn.net/Articles/164121/>
- libhugetlbfs
 - allows applications to use the huge page support in Linux with no need for source code modification
- kernel-headers package
 - replaces the `glibc-kernheaders` package
 - provides better suitability with the new `headers_install` feature of the 2.6.18 kernel
 - notable kernel header-related changes:
 - removed `<linux/compiler.h>` header file, as it is no longer useful
 - removed `__syscallX()` macros; user-space should use `syscall()` from the C library instead
 - removed `<asm/atomic.h>` and `<asm/bitops.h>` header files; C compiler provides its own atomic built-in functions better suitable for user-space programs
 - content previously protected with `#ifdef __KERNEL__` is now removed completely with the `unifdef` tool; defining `__KERNEL__` in order to view parts which should not be visible to user-space is no longer effective

- removed the `PAGE_SIZE` macro from some architectures, due to variance in page sizes; user-space should be using `sysconf (_SC_PAGE_SIZE)` or `getpagesize()`
- to provide better suitability for user-space, removed several header files and header content

Generic Feature Additions

- `kexec` and `kdump` (2.6.13)
 - `diskdump` and `netdump` have been replaced by `kexec` and `kdump`, which ensure faster boot-up and creation of reliable kernel vmcores for diagnostic purposes. For more information and configuration instructions, please refer to `/usr/share/doc/kexec-tools-<version>/kexec-kdump-howto.txt` (replace `<version>` with the corresponding version of the `kexec-tools` package installed).
 - Note that at present, virtualized kernels cannot use the `kdump` function.
- `inotify` (2.6.13)
 - user interface for this is through the following syscalls: `sys_inotify_init`, `sys_inotify_add_watch`, and `sys_inotify_rm_watch`.
- Process Events Connector (2.6.15)
 - reports `fork`, `exec`, id change, and exit events for all processes to user-space.
 - Applications that may find these events useful include accounting / auditing (for example, ELSA), system activity monitoring (for example, `top`), security, and resource management (for example, CKRM). Semantics provide the building blocks for features like per-user-namespace, "files as directories" and versioned file systems.
- Generic RTC (RealTime Clock) subsystem (2.6.17)
- `splice` (2.6.17)
 - new IO mechanism which avoids data copies when transferring data between applications
 - ref: <http://lwn.net/Articles/178199/>

File System / LVM

- EXT3
 - support for Extended Attributes in the body of large inode in ext3: saves space and improves performance in some cases (2.6.11)
- Device mapper multipath support
- ACL support for NFSv3 and NFSv4 (2.6.13)
- NFS: supports large reads and writes on the wire (2.6.16)
 - The Linux NFS client now supports transfer sizes of up to 1MB.

- VFS changes
 - The "shared subtree" patches have been merged. (2.6.15)
 - ref: <http://lwn.net/Articles/159077/>
- Big CIFS update (2.6.15)
 - features several performance improvements as well as support for Kerberos and CIFS ACL
- autofs4: updated to provide direct mount support for user-space autofs (2.6.18)
- cachefs core enablers (2.6.18)

Security


- Multilevel security implementation for SELinux (2.6.12)
- Audit subsystem
 - support for process-context based filtering (2.6.17)
 - more filter rule comparators (2.6.17)
- TCP/UDP getpeercon: enabled security-aware applications to retrieve the entire security context of a process on the other side of a socket using an IPSec security association. If only MLS-level information is needed or interoperability with legacy unix system is required, NetLabel can be used in place of IPSec.

Networking

- Added several TCP congestion modules (2.6.13)
- IPv6: supports several new sockopt / ancillary data in Advanced API (2.6.14)
- IPv4/IPv6: UFO (UDP Fragmentation Offload) Scatter-gather approach (2.6.15)
 - UFO is a feature wherein the Linux kernel network stack will offload the IP fragmentation functionality of large UDP datagram to hardware. This will reduce the overhead of stack in fragmenting the large UDP datagram to MTU-sized packets.
- Added nf_contrack subsystem (2.6.15)
 - The existing connection tracking subsystem in netfilter can only handle ipv4. There were two choices present to add connection tracking support for ipv6; either duplicate all of the ipv4 connection tracking code into an ipv6 counterpart, or (the choice taken by these patches) design a generic layer that could handle both ipv4 and ipv6 and thus requiring only one sub-protocol (TCP, UDP, etc.) connection tracking helper module to be written. In fact, nf_contrack is capable of working with any layer 3 protocol.
- IPV6
 - RFC 3484-compliant source address selection (2.6.15)
 - added support for Router Preference (RFC4191) (2.6.17)
 - added Router Reachability Probing (RFC4191) (2.6.17)

- added support for Multiple Routing Tables and Policy Routing
- Wireless updates
 - hardware crypto and fragmentation offload support
 - QoS (WME) support, "wireless spy support"
 - mixed PTK/GTK
 - CCMP/TKIP support and WE-19 HostAP support
 - BCM43xx wireless driver
 - ZD1211 wireless driver
 - WE-20, version 20 of the Wireless Extensions (2.6.17)
 - added the hardware-independent software MAC layer, "Soft MAC" (2.6.17)
 - added LEAP authentication type
- Added generic segmentation offload (GSO) (2.6.18)
 - can improve performance in some cases, though it needs to be enabled through `ethtool`
- DCCPv6 (2.6.16)

Added Hardware Support



NOTE

This section only enumerates the most generic features among many.

- x86-64 clustered APIC support (2.6.10)
- Infiniband support (2.6.11)
- Hot plug
 - added generic memory add/remove and supporting functions for memory hotplug (2.6.15)
- SATA/libata enhancements, additional hardware support
 - A completely reworked libata error handler; the result of all this work should be a more robust SATA subsystem which can recover from a wider range of errors.
 - Native Command Queuing (NCQ), the SATA version of tagged command queuing - the ability to have several I/O requests to the same drive outstanding at the same time. (2.6.18)
 - Hotplug support (2.6.18)
- EDAC support (2.6.16)

- The EDAC goal is to detect and report errors that occur within the system.
- Added a new ioatdma driver for the Intel(R) I/OAT DMA engine (2.6.18)

NUMA (Non-Uniform Memory Access) / Multi-core

- Cpusets (2.6.12)
 - Cpusets now provide a mechanism for assigning a set of CPUs and Memory Nodes to a set of tasks. Cpusets constrain the CPU and memory placement of tasks only to the resources within a task's current cpuset. These are essential in managing dynamic job placement on large systems.
- NUMA-aware slab allocator (2.6.14)
 - This creates slabs on multiple nodes and manages slabs in such a way that locality of allocations is optimized. Each node has its own list of partial, free and full slabs. All object allocations for a node occur from node-specific slab lists.
- Swap migration (2.6.16)
 - Swap migration allows the moving of physical location of pages between nodes in a NUMA system while the process is running.
- Huge pages (2.6.16)
 - Added NUMA policy support for huge pages: the `huge_zone_list()` function in the memory policy layer provides a list of zones ordered by NUMA distance. The `hugetlb` layer will walk that list looking for a zone that has available huge pages but is also in the `nodeset` of the current cpuset.
 - Huge pages now obey cpusets.
- Per-zone VM counters
 - provide zone-based VM statistics, which are necessary in determining what state of memory a zone is in
- Netfilter `ip_tables`: NUMA-aware allocation. (2.6.16)
- Multi-core
 - Added a new scheduler domain for representing multi-core with shared caches between cores. This makes it possible to make smarter cpu scheduling decisions on such systems, improving performance greatly for some cases. (2.6.17)
 - Power saving policy for the CPU scheduler: with `multicore/smt` cpus, the power consumption can be improved by leaving some packages idle while others do all the work, instead of spreading the tasks over all CPUs.

5. RELEASE NOTES FOR S390X

5.1. Installation-Related Notes

The following section includes information specific to the installation of Red Hat Enterprise Linux and the Anaconda installation program.



NOTE

In order to upgrade an already-installed Red Hat Enterprise Linux, you must use Red Hat Network to update those packages that have changed.

You may use Anaconda to perform a fresh installation of Red Hat Enterprise Linux 5 or to perform an upgrade from *the latest updated version* of Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5.

If you are copying the contents of the Red Hat Enterprise Linux 5 CD-ROMs (in preparation for a network-based installation, for example) be sure to copy the CD-ROMs for the operating system *only*. Do not copy the Supplementary CD-ROM, or any of the layered product CD-ROMs, as this will overwrite files necessary for Anaconda's proper operation. These CD-ROMs must be installed *after* Red Hat Enterprise Linux has been installed.

5.1.1. ISO Contents and Registration

The organization of software component packages into product specific-variants has changed from previous versions of Red Hat Enterprise Linux. The number of variants and ISO images has been reduced to only one: the Red Hat Enterprise Linux 5 Server.

In the future, additional repositories for a number of options that provide added functionality over the core Server distribution may be provided.

With optional content in the same tree or ISO image, it is important to avoid a mismatch between the components offered for installation and those covered by the subscription. Such a mismatch could result in an increased exposure to bug and vulnerability risks.

In order to ensure that the components offered are in sync with the subscription, Red Hat Enterprise Linux 5 requires entering an *Installation Number* that will be used to configure the installer to offer the right package set. This Installation Number is included in your subscription.

If you skip entering the Installation Number, this will result in a core Server or Desktop installation. Additional functionality can then be added manually at a later time. For more information about Installation Numbers, please refer to <http://www.redhat.com/apps/support/in.html>.

The Installation Number used during the installation process will be saved in `/etc/sysconfig/rhn/install-num`. When registering with Red Hat Network, this file will be referenced by `rhn_register` to automatically determine which appropriate child channels the system should be subscribed to.

5.1.2. New RPM GPG Signing Keys

A new release signing key is used to sign Red Hat Enterprise Linux 5 packages. When updating a system for the first time, you will be prompted to allow this key to be installed.

Signing keys are distributed in the following files:

- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release` – contains the public key for the new release signing key
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-auxiliary` – contains the public key for an auxiliary release signing key, currently not in use
- `/etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-former` – contains the public key for the previous release signing key, used for past Red Hat Enterprise Linux releases

5.1.3. Subversion

In Red Hat Enterprise Linux 5, the Subversion version control system is linked against Berkeley DB 4.3. If you are upgrading from Red Hat Enterprise Linux 4 and any Subversion repositories which use the Berkeley DB backend "BDB" (rather than the pure file system-based "FSFS" backend) have been created on the system, special care must be taken to ensure the repositories can be accessible after the upgrade. The following process must be performed on the Red Hat Enterprise Linux 4 system, *prior* to upgrading to Red Hat Enterprise Linux 5:

1. Shut down any running processes and ensure that no processes can access the repository (for example, `httpd`, `svnserve` or any local users with direct access).
2. Create a backup of the repository using the following command:

```
svnadmin dump /path/to/repository | gzip > repository-backup.gz
```

3. Run the `svnadmin recover` command on the repository:

```
svnadmin recover /path/to/repository
```

4. Delete any unused log files in the repository:

```
svnadmin list-unused-dblogs /path/to/repository | xargs rm -vf
```

5. Delete any remaining shared-memory files in the repository:

```
rm -f /path/to/repository/db/__.db.0*
```

5.1.4. Other Installation Notes

- If IDE/PATA (*Parallel ATA*) devices are configured in "100% Native" mode, some BIOSes may prevent the Red Hat Enterprise Linux 5 installation process from completing successfully. To prevent this from occurring, configure the IDE/PATA mode as "Legacy" in the BIOS.
- The *IBM System z* does not provide a traditional Unix-style physical console. As such, Red Hat Enterprise Linux 5 for the *IBM System z* does not support the `firstboot` functionality during initial program load.

To properly initialize setup for Red Hat Enterprise Linux 5 on the *IBM System z*, run the following commands after installation:

- `/usr/bin/setup` – provided by the `setuptools` package
- `/usr/bin/rhn_register` – provided by the `rhn-setup` package
- Red Hat Enterprise Linux 5 now uses the Open Source driver for the IBM 3590 tape drive. The kernel contains the driver.

```
/lib/modules/2.6.17-1.2519.4.26.el5/kernel/drivers/s390/char/tape_3590.ko
```

- When booting Anaconda with PXE using the parameter `ksdevice=bootif`, you will still be prompted for the ethernet interface to use during installation. If only one ethernet device is plugged in, use the parameter `ksdevice=link` instead. Alternatively, you can also specify the

interface manually.

5.2. Technology Previews

Technology Preview features are currently *not* supported under Red Hat Enterprise Linux 5 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. Erratas 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.

Stateless Linux

Included in this release of Red Hat Enterprise Linux 5 are enabling infrastructure pieces for *Stateless Linux*. Stateless Linux is a new way of thinking about how a system is to be run and managed, designed to simplify provisioning and management of large numbers of systems by making them easily replaceable. This is accomplished primarily by establishing prepared system images which get replicated and managed across a large number of stateless systems, running the operating system in a read-only manner (please refer to `/etc/sysconfig/readonly-root` for more details).

In its current state of development, the Stateless features are subsets of the intended goals. As such, the capability is being labeled as a technology preview.

The following is a list of the initial capabilities included in Red Hat Enterprise Linux 5:

- running a stateless image over NFS
- running a stateless image via loopback over NFS
- running on iSCSI

It is highly recommended that those interested in testing stateless code read the HOWTO at <http://fedoraproject.org/wiki/StatelessLinuxHOWTO> and join stateless-list@redhat.com.

GFS2

GFS2 is an evolutionary advancement based on the *GFS* file system. While fully functional, *GFS2* is not yet considered production-ready. *GFS2* is targeted to become fully supported in a subsequent Red Hat Enterprise Linux 5 update. There is also an in-place conversion utility, `gfs2_convert`, which can update the metadata of the older *GFS* file system format, converting it to a *GFS2* file system.

FS-Cache

FS-Cache is a local caching facility for remote file systems that allows users to cache NFS data on a locally mounted disk. To set up the *FS-Cache* facility, install the `cachefilesd` RPM and refer to the instructions in `/usr/share/doc/cachefilesd-<version>/README`.

Replace `<version>` with the corresponding version of the `cachefilesd` package installed.

Compiz

Compiz is an OpenGL-based compositing window manager. In addition to regular window management, **compiz** also acts as a compositing manager, coordinating and synchronizing the overall desktop redrawing to provide a smoother desktop experience with less flicker.

Compiz uses 3D hardware acceleration to render effects such as live thumbnail windows, window drop shadows, animated window minimizing and transitions between virtual desktops.

Due to limitations in the current rendering architecture, **compiz** cannot work correctly with direct rendering OpenGL applications or applications using the Xv extension. Such applications will exhibit harmless rendering artifacts; because of this, **compiz** is currently a technology preview.

Enhancement for Ext3

In Red Hat Enterprise Linux 5, the EXT3 file system capacity has been extended beyond 8TB to a maximum of 16TB. This capability is being included as a technology preview, and is targeted for full support in a future release of Red Hat Enterprise Linux 5.

AIGLX

AIGLX is a technology preview feature of the otherwise fully supported X server. It aims to enable GL-accelerated effects on a standard desktop. The project consists of the following:

- a lightly modified X server
- an updated Mesa package that adds new protocol support

By installing these components, you can have GL-accelerated effects on your desktop with very few changes, as well as the ability to enable and disable them at will without replacing your X server. *AIGLX* also enables remote GLX applications to take advantage of hardware GLX acceleration.

Frysk GUI

The goal of the *frysk* project is to create an intelligent, distributed, always-on system monitoring and debugging tool that allows developers and system administrators to:

- monitor running processes and threads (including creation and destruction events)
- monitor the use of locking primitives
- expose deadlocks
- gather data
- debug any given process by choosing it from a list or allowing **frysk** to open a source code (or other) window on a process that is crashing or misbehaving

In Red Hat Enterprise Linux 5 the **frysk** graphical user interface is a technology preview, whereas the **frysk** command line interface is fully supported.

Systemtap

Systemtap provides free software (GPL) infrastructure to simplify the gathering of information about the running Linux system. This assists diagnosis of a performance or functional problem. With the help of **systemtap**, developers no longer need to go through the tedious and disruptive instrument, recompile, install, and reboot sequence that may be otherwise required to collect data.

Dogtail

Dogtail is a GUI test tool and automation framework written in Python that uses *Accessibility* technologies to communicate with desktop applications.

Support for Indic Languages and Sinhalese

Red Hat Enterprise Linux 5 also features support for the following languages as technology preview:

- *Assamese*
- *Kannada*
- *Sinhalese*
- *Telugu*

For more information about how to install and enable support for these languages, refer to the *Internationalization* section of this document.

Installing to dm-multipath Devices

Anaconda now has the capability to detect, create, and install to *dm-multipath* devices. To enable this feature, add the parameter `mpath` to the kernel boot line.

Note that the parameter `mpath` may cause a boot failure if a device's `major:minor` number changes. This issue will be addressed in a future update of Red Hat Enterprise Linux 5.

Installation / Boot for iSCSI software initiator (open-iscsi)

Anaconda now provides the ability to install to an iSCSI device. Booting and installing is fully supported with the *QLogic qla4xxx* hardware initiator. However, the capability to install to an iSCSI device for the `open-iscsi` software initiator is currently considered a Technology Preview, due to the following issues:

- Text mode installation does not complete. You must do a graphical install, or an automated kickstart install.
- Media-based installations do not complete. You must do a network-based install.
- Depending on the timing of events, Anaconda may be unable to detect all the iSCSI targets or LUNs. When this occurs, use the installer shell to configure the storage through iSCSI commands.
- The `iscsid` daemon may not properly start. Such an occurrence will prevent the system from handling all iSCSI errors, such as network problems, SCSI/iSCSI timeouts, and target errors. To confirm that the `iscsid` daemon is running, run the command `iscsiadm -m session -i` and check that the line, `Internal iscsid Session State:` reports a value (it can be any value).
- On certain iSCSI target implementations, the system may hang during shutdown.
- On certain iSCSI target implementations, the system may hang during reboot. To avoid this, shutdown the system and boot it up again (instead of rebooting directly from a session).
- Booting from iSCSI devices on the *IBM System p* does not work reliably. While installation on an iSCSI device may appear to succeed, the resulting installation will not boot properly.
- On the first boot after install, you may receive SELinux errors such as the following:

```
kernel: audit(1169664832.270:4): avc: denied { read
} for pid=1964 comm="iscsid"
```

To work around this, boot the system with the kernel parameter `enforcing=0`. Once the system has properly booted, use the command `setenforce 1` to restore enforcing mode.

These limitations will be addressed in a future Red Hat Enterprise Linux 5 update.

5.3. Known Issues

- Host bus adapters that use the MegaRAID driver must be set to operate in "Mass Storage" emulation mode, not in "I2O" emulation mode. To do this, perform the following steps:

1. Enter the MegaRAID BIOS **Set Up Utility**.
2. Enter the **Adapter settings** menu.
3. Under **Other Adapter Options**, select **Emulation** and set it to **Mass Storage**.

If the adapter is incorrectly set to "I2O" emulation, the system will attempt to load the i2o driver. This will fail, and prevent the proper driver from being loaded.

Previous Red Hat Enterprise Linux releases generally do not attempt to load the I2O driver before the MegaRAID driver. Regardless of this, the hardware should always be set to "Mass Storage" emulation mode when used with Linux.

- Red Hat Enterprise Linux 5 includes `openmpi-1.1.1-4.el5` (from the OFED 1.1 distribution), which has been discovered to eventually quit working entirely. This happens after the `openmpi` stack works as expected for a varying amount of time.

For updated versions of `openmpi`, please check <http://people.redhat.com/dledford/Infiniband/openmpi>

- Installing Windows Server 2003 as a guest on a fully virtualized Red Hat Enterprise Linux 5 system ends unexpectedly after completing the first stage of installation. When this happens, the graphical console window closes, and the guest disappears from the **Virtual Machine Manager's** list of machines, resulting in a **Broken pipe** error.

This issue will be resolved in an upcoming Red Hat Enterprise Linux 5 update. To work around this, use the following command at the terminal:

```
xm create /etc/xen/<name of guest machine>
```

Afterwards, open the virtual machine.

- When attempting to create a fully virtualized Windows Server 2003 from a CD / DVD, the second stage of the guest install will not continue upon reboot.

To work around this, edit `/etc/xen/<name of guest machine>` by properly appending an entry for the CD / DVD device.

If an installation to a simple file is used as a virtual device, the `disk` line of `/etc/xen/<name of guest machine>` will read like the following:

```
disk = [ 'file:/PATH-OF-SIMPLE-FILE,hda,w' ]
```

A DVD-ROM device located on the host as `/dev/dvd` can be made available to stage 2 of the installation as `hdc` by appending an entry like `'phy:/dev/dvd,hdc:cdrom,r'`. As such, the disk line should now read as follows:

```
disk = [ 'file:/opt/win2003-sp1-20061107,hda,w',
         'phy:/dev/dvd,hdc:cdrom,r' ]
```

The precise device path to use may vary depending on your hardware.

- `rmmod xennet` causes `domU` to crash; this is caused by a grant table issue in the Virtualization feature. Due to the current inability of the Virtualization feature to asynchronously release grant table operations, it is unsafe to unload the `xennet` module in guests. In such situations, grant tables are used to perform backend-frontend communication, and there is no guarantee that the backend will release the references, leading to an inevitable memory leak.

This issue will be resolved in the next minor release of Red Hat Enterprise Linux 5. At present, users are advised *not* to unload the `xennet` module in guests.

- Running `ethtool eth0` outputs incomplete information about the ethernet card settings. This only occurs in systems running a virtualized kernel, since the Virtualization feature uses a networking setup where the physical ethernet device is identified as `peth0`. As such, the correct command for retrieving information about the physical ethernet device is `ethtool peth0`.
- When using Red Hat Enterprise Linux 5 on a machine with an *nVidia CK804* chipset installed, you may receive kernel messages similar to the following:

```
kernel: assign_interrupt_mode Found MSI capability
kernel: pcie_portdrv_probe->Dev[005d:10de] has invalid IRQ. Check
vendor BIOS
```

These messages indicate that certain PCI-E ports are not requesting IRQs. Further, these messages do not, in any way, affect the operation of the machine.

- Some *Cisco Aironet Wireless* devices prevent **NetworkManager** from storing connection details for wireless networks that do not broadcast an SSID. This is caused by a *Cisco Aironet Wireless* device firmware limitation.
- Laptops that have the *Cisco Aironet MPI-350* wireless card equipped may hang trying to get a DHCP address during any network-based installation using the wired ethernet port.

To work around this, use local media for your installation. Alternatively, you can disable the wireless card in the laptop BIOS prior to installation (you can re-enable the wireless card after completing the installation).

- Currently, `system-config-kickstart` does not support package selection and deselection. When using `system-config-kickstart`, the **Package Selection** option indicates that it is disabled. This is because `system-config-kickstart` uses `yum` to gather group information, but is unable to configure `yum` to connect to Red Hat Network.

This issue is currently being investigated for resolution by the next minor release of Red Hat Enterprise Linux 5. At present, you need to update package sections in your kickstart files manually. When using `system-config-kickstart` to open a kickstart file, it will preserve all

package information in it and write it back out when you save.

- 4-socket AMD *Sun Blade X8400 Server Modules* systems that do not have memory configured in `node 0` will panic during boot. Systems should be configured with memory in `node 0` to prevent the kernel panic.
- Installing to LVM mirror devices through Anaconda is currently not supported. This capability will be added in a future update of Red Hat Enterprise Linux 5.
- When installing Red Hat Enterprise Linux 5 from a directory on an NFS server containing Red Hat Enterprise Linux ISO images, Anaconda may display the following error message:

```
Unable to read package metadata. This may be due to a missing
repdata directory.
Please ensure that your install tree has been correctly generated.
Cannot
open/read repomd.xml file for repository:
```

This problem occurs if the directory holding the ISO images also contains a partially unpacked installation tree (for example, the `/images` directory from the first ISO). The presence of such directories results in the error stated above.

To prevent this error, unpack trees only to directories other than the one containing the installation ISO images.

- Boot-time logging to `/var/log/boot.log` is not available in this release of Red Hat Enterprise Linux 5. An equivalent functionality will be added in a future update of Red Hat Enterprise Linux 5.
- Neither `kexec` nor `kdump` are able to dump onto disks attached to an `accraid` controller.

To work around this issue, use `scp` for network dumping. Alternatively, you can also dump onto a disk through a different controller.

- If a *Hipersocket* (iQDIO) network interface is being used on a Red Hat Enterprise Linux 5 system, shutting down may cause the system to hang while trying to unload the `iptables` modules. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.
- Installing a fully virtualized guest using split installation media -- specifically, multiple CD-ROMs -- may fail when required to switch between installation CDs. During the guest OS installation process, users may be prevented from mounting or ejecting installation CDs, which prevents the installation from completing.

As such, it is recommended that you use the QEMU monitor console to switch CD-ROM images during the guest OS installation process. The procedure is as follows:

1. Open a graphical VNC console to the guest OS.
2. Unmount the CD-ROM device in the guest OS.
3. Switch to the QEMU monitor console by pressing `Ctrl-Alt-2`.
4. Run the command `eject hdc`.
5. Run the command `change hdc <path to the CD-ROM in host system>`.

6. Switch back to the guest OS console by pressing **Ctrl-Alt-1**.

7. Mount the CD-ROM device in the guest OS.

Note that when using a regular VNC client the host X server may encounter some difficulty interpreting the **Ctrl-Alt-2** and **Ctrl-Alt-1** command. To work around this in `virt-manager`, use *sticky keys*. Pressing **Ctrl** three times makes it "sticky" until the next non-modifier is pressed. As such, to send **Ctrl-Alt-1**, press **Ctrl** twice before pressing **Ctrl-Alt-1**.

- The Red Hat Enterprise Linux 5 *Driver Update Model* creates modified `initrd` images whenever a `kmod` package that includes a bootpath-modifying driver is installed. In time, the number of backup `initrd` images may soon fill the `/boot` partition, particularly if the system undergoes a sizeable number of driver updates.

As such, it is recommended that you monitor the free space on the `/boot` partition if you regularly perform driver updates. You can free up more space in `/boot` by removing older `initrd` images; these files end in `.img0`, `.img1`, `.img2`, and so on.

- Autorun on removable media is currently disabled. To install packages from the Red Hat Enterprise Linux Supplementary CD, launch the CD installer manually using the following command:

```
system-cdinstall-helper /media/path-to-mounted-drive
```

- When upgrading from Red Hat Enterprise Linux 4 to Red Hat Enterprise Linux 5, the Deployment Guide is not automatically installed. You need to use `pirut` to manually install it after completing the upgrade.
- An `autofs` bug prevents multi-mounts from working properly.

During an expiry, if the last multi-mount component to be checked does not have a mount associated with it while other components are busy, `autofs` erroneously determines the multi-mount to be expirable. This causes the multi-mount to be partially expired, resulting in the multi-mount becoming unresponsive to further mount requests and expire runs.

To permanently resolve this problem, update `autofs` using the command `yum update autofs`.

- The system may not successfully reboot into a `kexec/kdump` kernel if X is running and using a driver other than `vesa`. This problem only exists with *ATI Rage XL* graphics chipsets.

If X is running on a system equipped with *ATI Rage XL*, ensure that it is using the `vesa` driver in order to successfully reboot into a `kexec/kdump` kernel.

5.4. General Information

This section contains general information not specific to any other section of this document.

Red Hat Enterprise Linux Deployment Guide

This release of Red Hat Enterprise Linux includes a fairly comprehensive Deployment Guide. To access it, go to **System** (on the top panel) => **Documentation** => **Red Hat Enterprise Linux Deployment Guide**.

It is the intention of Red Hat to provide fully localized versions of the Deployment Guide for all supported languages. If you have installed a localized version of the Deployment Guide, it is recommended that you update it when a new version becomes available through Red Hat Network.

Web Server Packaging Changes

Red Hat Enterprise Linux 5 now includes version 2.2 of the Apache HTTP Server. This release brings a number of improvements over the 2.0 series, including:

- improved caching modules (`mod_cache`, `mod_disk_cache`, `mod_mem_cache`)
- a new structure for authentication and authorization support, replacing the authentication modules provided in previous versions
- support for proxy load balancing (`mod_proxy_balancer`)
- support for handling large files (namely, greater than 2GB) on 32-bit platforms

The following changes have been made to the default `httpd` configuration:

- The `mod_cern_meta` and `mod_asis` modules are no longer loaded by default.
- The `mod_ext_filter` module is now loaded by default.

If you are upgrading from a previous release of Red Hat Enterprise Linux, the `httpd` configuration will need to be updated for `httpd 2.2`. For more information, refer to <http://httpd.apache.org/docs/2.2/upgrading.html>.

Note that any third-party modules compiled for `httpd 2.0` must be rebuilt for `httpd 2.2`.

php

Version 5.1 of PHP is now included in Red Hat Enterprise Linux 5, which includes a number of changes to the language along with significant performance improvements. Some scripts might need to be edited for use with the new version; please refer to the link below for more information on migrating from PHP 4.3 to PHP 5.1:

<http://www.php.net/manual/en/migration5.php>

The `/usr/bin/php` executable is now built using the CLI command-line SAPI, rather than the CGI SAPI. Use `/usr/bin/php-cgi` for CGI SAPI. The `php-cgi` executable also includes FastCGI support.

The following extension modules have been added:

- the `mysqli` extension, a new interface designed specifically for MySQL 4.1 (included in the `php-mysql` package)
- `date`, `hash`, `Reflection`, `SPL` and `SimpleXML` (built-in with the `php` package)
- `pdo` and `pdo_sqlite` (in the `php-pdo` package)
- `pdo_mysql` (in the `php-mysql` package)
- `pdo_pgsql` (in the `php-pgsql` package)
- `pdo_odbc` (in the `php-odbc` package)

- `soap` (in the `php-soap` package)
- `xmlreader` and `xmlwriter` (in the `php-xml` package)
- `dom` (replacing the `domxml` extension in the `php-xml` package)

The following extension modules are no longer included:

- `dbx`
- `dio`
- `yp`
- `overload`
- `domxml`

The PEAR Framework

The PEAR framework is now packaged in the `php-pear` package. Only the following PEAR components are included in Red Hat Enterprise Linux 5:

- `Archive_Tar`
- `Console_Getopt`
- `XML_RPC`

Encrypted Swap Partitions and Non-root File Systems

Red Hat Enterprise Linux 5 now provides basic support for encrypted swap partitions and non-root file systems. To use these features, add the appropriate entries to `/etc/crypttab` and reference the created devices in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry:

```
my_swap /dev/hdb1 /dev/urandom swap,cipher=aes-cbc-essiv:sha256
```

This creates the encrypted block device `/dev/mapper/my_swap`, which can be referenced in `/etc/fstab`.

Below is a sample `/etc/crypttab` entry for a file system volume:

```
my_volume /dev/hda5 /etc/volume_key cipher=aes-cbc-essiv:sha256
```

The `/etc/volume_key` file contains a plaintext encryption key. You can also specify `none` as the key file name; this configures the system to ask for the encryption key during boot instead.

It is recommended to use LUKS (Linux Unified Key Setup) for setting up file system volumes. To do this, follow these steps:

1. Create the encrypted volume using `cryptsetup luksFormat`.
2. Add the necessary entry to `/etc/crypttab`.

3. Set up the volume manually using `cryptsetup luksOpen` (or reboot).
4. Create a file system on the encrypted volume.
5. Add the necessary entry to `/etc/fstab`.

mount and umount

The `mount` and `umount` commands no longer directly support NFS; a built-in NFS client no longer exists. A separate `nfs-utils` package, which provides `/sbin/mount.nfs` and `/sbin/umount.nfs` helpers, must be installed for this.

CUPS Printer Browsing

CUPS printer browsing over a local subnet can be configured using the graphical tool `system-config-printer`. It can also be done using the CUPS web interface, <http://localhost:631/>.

To use directed broadcasts for printer browsing between subnets, open `/etc/cups/cupsd.conf` on the clients and replace `BrowseAllow @LOCAL` with `BrowseAllow ALL`.

ATI and R500 Support

ATI graphics cards based on the R500 chipset are supported for the `vesa` driver only, and are not supported by Red Hat Enterprise Linux 5 on external monitors, LCD projectors or accelerated 3D support.

up2date and yum

`up2date` is being deprecated in favor of `yum` (Yellowdog Updater Modified). As such, it is advisable that you revise any `up2date`-dependent scripts your system is using accordingly. For more information about `yum`, consult its man page with the command `man yum`; you can also consult the installed documentation under the directories `/usr/share/doc/yum-<version>` and `/usr/share/doc/yum-metadata-parser-<version>` (replace `<version>` with the corresponding version of `yum` and `yum-metadata-parser` installed).

OpenLDAP Server and Red Hat Directory Server

Red Hat Directory Server is an LDAP-based server that centralizes enterprise and network data into an OS-independent, network-based registry. It is set to replace OpenLDAP server components, which will be deprecated *after* Red Hat Enterprise Linux 5. For more information about Red Hat Directory Server, refer to <http://www.redhat.com/software/rha/directory/>.

IBM System z Support for LCS and QDIO-enabled Devices

Red Hat Enterprise Linux 5 for the *IBM System z* supports LCS (*LAN Channel Station*) and QDIO-enabled (*Queued Direct Input/Output*) devices. The CTC (channel-to-channel) and NETIUCV (*Inter-user Communication Vehicle Network Device Support*) drivers are no longer supported.

i810 Driver and i830 Support

The *i810* driver supports all integrated Intel graphics chipsets, from *i810* to *i965*. However, the support for *i830* (and newer) chipsets is limited; the *i810* driver can only set modes listed in the video BIOS. If your machine has an *i830* or newer chipset installed, run the following command to determine what the available modes are:

```
grep Mode: /var/log/Xorg.0.log
```

Modes marked with an asterisk (*) are available for selection.

Many laptop video BIOSes do not supply a mode that matches the native panel size. Therefore the chosen mode may appear stretched, distorted, or with black borders. As such, if your chosen mode does not display properly, you need a BIOS update from your hardware vendor for the native panel size to work correctly.

Intel PRO/Wireless 3945ABG Network Connection Support

This release of Red Hat Enterprise Linux 5 includes support for the *ipw3945* (*Intel PRO/Wireless 3945ABG Network Connection*) adapter. The Red Hat Enterprise Linux 5 Supplementary disc contains the driver, regulatory daemon and firmware needed to support this adapter.

To enable support for the ipw3945 wireless adapter, search the Red Hat Enterprise Linux 5 Supplementary disc for packages with filenames containing "3945" and install them.

rawio

rawio is a deprecated interface; however, Red Hat Enterprise Linux 5 still includes support for it. If you have an application that performs device access using *rawio*, it is highly recommended that you modify your application to open the block device with the `O_DIRECT` flag. The *rawio* interface will remain throughout the life of Red Hat Enterprise Linux 5, but is a candidate for removal in a future release.

Currently, AIO (*Asynchronous I/O*) on file systems is only supported in `O_DIRECT` or non-buffered mode. Further, note that the asynchronous poll interface is no longer present, and that AIO on pipes is no longer supported.

ctmpc

ctmpc is a deprecated driver; however, it will still be included throughout the life of Red Hat Enterprise Linux 5. Note that it is a candidate for removal from future releases.

Policy Modules and semanage Support

Red Hat Enterprise Linux 5 now supports *policy modules* and *semanage*. Policy modules simplify the creation and distribution of policy customizations and third-party policies through the use of the `semodule` and `checkmodule` tools.

The *semanage* tool is a policy management tool that modifies the SELinux configuration. It also allows you to configure file contexts, networking component labeling, and user mappings for Linux-to-SELinux.

raw Device Mapping

The raw devices interface has been deprecated in Red Hat Enterprise Linux 5; raw device mapping is now configured via `udev` rules.

To configure raw device mapping, add the appropriate entries to `/etc/udev/rules.d/60-raw.rules` in the following formats:

- For device names:

```
ACTION=="add", KERNEL="device name", RUN+="raw /dev/raw/rawX %N"
```

- For major / minor numbers:

```
ACTION=="add", ENV{MAJOR}="A", ENV{MINOR}="B", RUN+="raw /dev/raw/rawX %M %m"
```

Replace `<device name>` with the name of the device you need to bind (for example, `/dev/sda1`). "A" and "B" are the major / minor numbers of the device you need to bind, and `X` is the raw device number that you want the system to use.

If you have a large, pre-existing `/etc/sysconfig/rawdevices` file, convert it with the following script:

```
#!/bin/sh

grep -v "^ *#" /etc/sysconfig/rawdevices | grep -v "^$" | while read dev
major minor ; do
  if [ -z "$minor" ]; then
    echo "ACTION==\"add\", KERNEL==\"${major##/dev/}\",
RUN+=\"/usr/bin/raw $dev %N\""
  else
    echo "ACTION==\"add\", ENV{MAJOR}==\"$major\", ENV{MINOR}==\"$minor\",
RUN+=\"/usr/bin/raw $dev %M %m\""
  fi
done
```

QLogic Support

Red Hat Enterprise Linux 5 supports the *QLogic* family of iSCSI HBA (*Host Bus Adapters*). At present, only the iSCSI interface to these boards are supported (using the `qla4xxx` driver).

In addition, Red Hat does not currently support these boards as Ethernet NIC, as this capability requires the `qla3xxx` driver. This issue will be addressed in an upcoming minor release of Red Hat Enterprise Linux 5.

SCSI Support

SCSI support over devices attached through fiber channel on the *IBM System z* has the following restrictions:

- Due to serviceability issues, it is not recommended to run a system with storage devices that are only attached through FCP (*Fiber Channel Protocol*). In order for the system to take stand-alone dumps, another storage device must be attached to the system through ESCON (*Enterprise System Connection*) or FICON (*Fiber Connectivity*).

Note that for large systems, the amount of data to be dumped may exceed the capacity available on the ESCON/FICON disk storage controller or tape drive. In such cases, the ESCON/FICON storage devices cannot be used for data dumps.

- The performance analysis for FCP-attached storage included in Red Hat Enterprise Linux 5 does not cover FCP-specific Linux functions on the *IBM System z*. These functions are required to obtain I/O performance profiles similar to those which can be obtained for ESCON/FICON storage devices. Hence, the capability to analyze and resolve FCP performance problems (as well as plan FCP I/O capacity) is limited.

As such, users are advised *not* to implement I/O performance-critical production workloads when running Red Hat Enterprise Linux 5 on the *IBM System z* with SCSI over devices attached through fiber channel.

- Red Hat Enterprise Linux 5 does not include a version of the HBA API library for Linux on the *IBM System z*. As such, this release does not support SAN management middleware that requires an HBA API library.

IBM System z Instruction Set

In order to optimally exploit the *IBM System z* instruction set for 31-bit applications, it is recommended that you use the `gcc` option `-march=z900`. For 64-bit applications, the `gcc` will exploit the *IBM System z* instruction set by default.

iSeries Access for Linux

The *iSeries ODBC Driver* for Linux has been replaced by the *iSeries Access for Linux*, which can be downloaded at the following link:

<http://www.ibm.com/eserver/series/access/linux/>

The *iSeries Access for Linux* offers Linux-based access to iSeries servers, and allows you to:

- Access the DB2 UDB (*Universal Database*) for iSeries using its ODBC Driver
- Establish a 5250 session to an iSeries server from a Linux client
- Access the DB2 UDB via the EDRS (*Extended Dynamic Remote SQL*) driver
- Support 32-bit (i386 and PowerPC) and 64-bit (x86-64 and PowerPC) platforms

IBM Power4 iSeries

Red Hat Enterprise Linux no longer supports the *IBM Power4 iSeries*.

5.5. Driver Update Program

This section includes information on the implementation of the Red Hat Enterprise Linux 5 Driver Update Program.

Kernel Module Packages

On Red Hat Enterprise Linux 5, it is possible to build updated kernel module packages that depend upon the current kernel ABI version and not on a specific kernel release number. This facilitates building kernel modules that can be used against a range of Red Hat Enterprise Linux 5 kernels, rather than a single release. The project website at <http://www.kerneldrivers.org/> contains more information about the packaging process, as well as several examples.

Note that the following issues have also been identified:

- A minor ABI compatibility issue affecting the network `rhe15_drivers_s390_net_ga` symbol grouping on s390x kernels has been identified. Symbols beginning with `iucv_` should not be used by third party drivers as they may be updated sometime throughout the life of Red Hat Enterprise Linux 5. Similarly, on ppc64, the structure `spu_priv1_ops` should not be used by third parties.
- Bootpath drivers distributed as `kmod` packages are not officially supported.
- Overriding existing in-kernel drivers are not currently supported.

These issues will be addressed in a future update of Red Hat Enterprise Linux 5.

Kernel Module Loading

The module loading behavior on Red Hat Enterprise Linux 5 has changed from previous releases of Red Hat Enterprise Linux. The modules shipped in the Red Hat Enterprise Linux 5 kernel package

are signed, as was the case in Red Hat Enterprise Linux 4. On Red Hat Enterprise Linux 5 kernels, however, it is no longer possible to load a signed module from another kernel build.

This means that a module shipped with the initial Red Hat Enterprise Linux 5 distribution cannot be loaded in future updated kernels. This helps prevent users from loading unsupported modules on a system. Red Hat only supports modules that are signed and included in a distribution.

If you want to load an older module, you can try rebuilding it without a signature. Alternatively, you can remove the signature from the binary file using the following command:

```
objcopy -R .module_sig <module name>-mod.ko <module name>-nosig.ko
```

It is recommended that you consult with a designated Red Hat Global Support Representative before attempting to load unsigned modules.

5.6. Internationalization

This section includes information on language support under Red Hat Enterprise Linux 5.

Input Methods

SCIM (Smart Common Input Method) has replaced IIMF as the input method system for Asian and other languages in this release. The default GTK Input Method Module for SCIM is provided by `scim-bridge`; in Qt, it is provided by `scim-qtim`.

Below are the default trigger hotkeys for different languages:

- All languages: **Ctrl-Space**
- Japanese: **Zenkaku-Hankaku** or **Alt-`**
- Korean: **Shift-Space**

If SCIM is installed, it runs by default for all users.

After installing or removing SCIM engine packages, it is recommended to start a new desktop session in order for the changes to be reflected in the SCIM language menu.

Language Installation

To enable additional language support for some Asian languages, you need to install the necessary language support packages. Below is a list of these languages and the command you need to run (as root) to install their corresponding language support packages:

- Assamese – `yum install fonts-bengali m17n-db-assamese scim-m17n`
- Bengali – `yum install fonts-bengali m17n-db-bengali scim-m17n`
- Chinese – `yum install fonts-chinese scim-chewing scim-pinyin scim-tables-chinese`
- Gujarati – `yum install fonts-gujarati m17n-db-gujarati scim-m17n`
- Hindi – `yum install fonts-hindi m17n-db-hindi scim-m17n`
- Japanese – `yum install fonts-japanese scim-anthy`

- `Kannada – yum install fonts-kannada m17n-db-kannada scim-m17n`
- `Korean – yum install fonts-korean scim-hangul`
- `Malayalam – yum install fonts-malayalam m17n-db-malayalam scim-m17n`
- `Marathi – yum install fonts-hindi m17n-db-marathi scim-m17n`
- `Oriya – yum install fonts-oriya m17n-db-oriya scim-m17n`
- `Punjabi – yum install fonts-punjabi m17n-db-punjabi scim-m17n`
- `Sinhala – yum install fonts-sinhala m17n-db-sinhala scim-m17n`
- `Tamil – yum install fonts-tamil m17n-db-tamil scim-m17n`
- `Telugu – yum install fonts-telugu m17n-db-telugu scim-m17n`

It is also recommended that you install `scim-bridge-gtk` and `scim-qtim` when enabling additional language support. The `scim-bridge-gtk` package prevents possible binary conflicts with third-party applications linked against older versions of `libstdc++`.

Note that additional language support packs are also available for **OpenOffice** (`openoffice.org-langpack-<language code>_<locale>`) and **KDE** (`kde-i18n-<language>`). These packages can also be installed through `yum`.

im-chooser

A new user configuration tool called `im-chooser` has been added, which allows you to easily disable or enable the usage of input methods on your desktop. So if SCIM is installed but you do not wish to run it on your desktop, you can disable it using `im-chooser`.

xinputrc

At X startup, `xinput.sh` now sources `~/xinputrc` or `/etc/X11/xinit/xinputrc` instead of searching config files under `~/xinput.d/` or `/etc/xinit/xinput.d/`.

Pango Support in Firefox

Firefox in Red Hat Enterprise Linux 5 is built with Pango, which provides better support for certain scripts, such as Indic and some CJK scripts.

To disable the use of Pango, set `MOZ_DISABLE_PANGO=1` in your environment before launching Firefox.

Fonts

Support is now available for synthetic emboldening of fonts that do not have a bold face.

New fonts for Chinese have been added: AR PL ShanHeiSun Uni (`uming.ttf`) and AR PL ZenKai Uni (`ukai.ttf`). The default font is AR PL ShanHeiSun Uni, which contains embedded bitmaps. If you prefer outline glyphs, add the following section to your `~/font.conf` file:

```
<fontconfig>
  <match target="font">
    <test name="family" compare="eq">
      <string>AR PL ShanHeiSun Uni</string>
```

```
</test>
<edit name="embeddedbitmap" mode="assign">
  <bool>>false</bool>
</edit>
</match>
</fontconfig>
```

gtk2 IM submenu

The Gtk2 context menu IM submenu no longer appears by default. You can enable it on the command line with the following command:

```
gconftool-2 --type bool --set
'/desktop/gnome/interface/show_input_method_menu' true
```

Support for text installation on CJK

CJK (Chinese, Japanese, and Korean) rendering support has been removed from the Anaconda text installation. The text installation method is being deprecated in the long term, as the GUI installation, VNC and kickstart methods are preferred.

gtk+ deprecation

The following packages are deprecated and scheduled for removal in Red Hat Enterprise Linux:

- **gtk+**
- **gdk-pixbuf**
- **glib**

These packages are being deprecated in favor of the **gtk2** stack, which offers better functionality particularly in terms of internationalization and font handling.

CJK input on console

If you need to display Chinese, Japanese, or Korean text on the console, you need to set up a framebuffer; afterwards, install **bogl-bterm**, and run **bterm** on the framebuffer.

5.7. Kernel Notes

This section notes the differences between 2.6.9 (on which Red Hat Enterprise Linux 4 is based) and 2.6.18 (which Red Hat Enterprise Linux 5 will inherit) as of July 12, 2006. Additional features which we are currently working on upstream (for example, virtualization) that will appear late in 2.6.18 or 2.6.19 are not highlighted here. In other words, this list only shows what is already included in the upstream Linux tree; not what is currently in development. Consequently, this list is not a final, or complete list of the new Red Hat Enterprise Linux 5 features, although it does give a good overview of what can be expected. Also, note that this section only picks out highlights of upstream changes, and as such it is not fully comprehensive. It does not include mention of several low-level hardware support enhancements and device driver info.

The following is a good source for a next level-of-detail view:

<http://kernelnewbies.org/LinuxChanges>

Performance / Scalability

- Big Kernel Lock preemption (2.6.10)
- Voluntary preemption patches (2.6.13) (subset in Red Hat Enterprise Linux 4)
- Lightweight user-space priority inheritance (PI) support for futexes, useful for real-time applications (2.6.18)
 - ref: <http://lwn.net/Articles/178253/>
- New 'mutex' locking primitive (2.6.16)
- High resolution timers (2.6.16)
 - In contrast to the low-resolution timeout API implemented in kernel/timer.c, hrtimers provide finer resolution and accuracy depending on system configuration and capabilities. These timers are currently used for itimers, POSIX timers, nanosleep and precise in-kernel timing.
- Modular, on-the-fly switchable I/O schedulers (2.6.10)
 - This was adjustable only by boot option in Red Hat Enterprise Linux 4 (also system-wide instead of per-queue).
- New Pipe implementation (2.6.11)
 - 30-90% performance improvement in pipe bandwidth
 - circular buffer allows more buffering than blocking writers
- "Big Kernel Semaphore": turns the Big Kernel Lock into a semaphore
 - reduces latency by breaking up long lock hold times and adding voluntary preemption
- X86 "SMP alternatives"
 - optimizes a single kernel image at runtime according to the available platform
 - ref: <http://lwn.net/Articles/164121/>
- kernel-headers package
 - replaces the `glibc-kernheaders` package
 - provides better suitability with the new `headers_install` feature of the 2.6.18 kernel
 - notable kernel header-related changes:
 - removed `<linux/compiler.h>` header file, as it is no longer useful
 - removed `__syscallX()` macros; user-space should use `syscall()` from the C library instead
 - removed `<asm/atomic.h>` and `<asm/bitops.h>` header files; C compiler provides its own atomic built-in functions better suitable for user-space programs
 - content previously protected with `#ifdef __KERNEL__` is now removed completely with the `unifdef` tool; defining `__KERNEL__` in order to view parts which should not be visible to user-space is no longer effective

- removed the `PAGE_SIZE` macro from some architectures, due to variance in page sizes; user-space should be using `sysconf (_SC_PAGE_SIZE)` or `getpagesize()`
- to provide better suitability for user-space, removed several header files and header content

Generic Feature Additions

- `kexec` and `kdump` (2.6.13)
 - `diskdump` and `netdump` have been replaced by `kexec` and `kdump`, which ensure faster boot-up and creation of reliable kernel vmcores for diagnostic purposes. For more information and configuration instructions, please refer to `/usr/share/doc/kexec-tools-<version>/kexec-kdump-howto.txt` (replace `<version>` with the corresponding version of the `kexec-tools` package installed).
 - Note that at present, virtualized kernels cannot use the `kdump` function.
- `inotify` (2.6.13)
 - user interface for this is through the following syscalls: `sys_inotify_init`, `sys_inotify_add_watch`, and `sys_inotify_rm_watch`.
- Process Events Connector (2.6.15)
 - reports `fork`, `exec`, id change, and exit events for all processes to user-space.
 - Applications that may find these events useful include accounting / auditing (for example, ELSA), system activity monitoring (for example, `top`), security, and resource management (for example, CKRM). Semantics provide the building blocks for features like per-user-namespace, "files as directories" and versioned file systems.
- Generic RTC (RealTime Clock) subsystem (2.6.17)
- `splice` (2.6.17)
 - new IO mechanism which avoids data copies when transferring data between applications
 - ref: <http://lwn.net/Articles/178199/>

File System / LVM

- EXT3
 - support for Extended Attributes in the body of large inode in ext3: saves space and improves performance in some cases (2.6.11)
- Device mapper multipath support
- ACL support for NFSv3 and NFSv4 (2.6.13)
- NFS: supports large reads and writes on the wire (2.6.16)
 - The Linux NFS client now supports transfer sizes of up to 1MB.

- VFS changes
 - The "shared subtree" patches have been merged. (2.6.15)
 - ref: <http://lwn.net/Articles/159077/>
- Big CIFS update (2.6.15)
 - features several performance improvements as well as support for Kerberos and CIFS ACL
- autofs4: updated to provide direct mount support for user-space autofs (2.6.18)
- cachefs core enablers (2.6.18)

Security

- Multilevel security implementation for SELinux (2.6.12)
- Audit subsystem
 - support for process-context based filtering (2.6.17)
 - more filter rule comparators (2.6.17)
- TCP/UDP getpeercon: enabled security-aware applications to retrieve the entire security context of a process on the other side of a socket using an IPSec security association. If only MLS-level information is needed or interoperability with legacy unix system is required, NetLabel can be used in place of IPSec.

Networking

- Added several TCP congestion modules (2.6.13)
- IPv6: supports several new sockopt / ancillary data in Advanced API (2.6.14)
- IPv4/IPv6: UFO (UDP Fragmentation Offload) Scatter-gather approach (2.6.15)
 - UFO is a feature wherein the Linux kernel network stack will offload the IP fragmentation functionality of large UDP datagram to hardware. This will reduce the overhead of stack in fragmenting the large UDP datagram to MTU-sized packets.
- Added nf_contrack subsystem (2.6.15)
 - The existing connection tracking subsystem in netfilter can only handle ipv4. There were two choices present to add connection tracking support for ipv6; either duplicate all of the ipv4 connection tracking code into an ipv6 counterpart, or (the choice taken by these patches) design a generic layer that could handle both ipv4 and ipv6 and thus requiring only one sub-protocol (TCP, UDP, etc.) connection tracking helper module to be written. In fact, nf_contrack is capable of working with any layer 3 protocol.
- IPV6
 - RFC 3484-compliant source address selection (2.6.15)
 - added support for Router Preference (RFC4191) (2.6.17)
 - added Router Reachability Probing (RFC4191) (2.6.17)

- added support for Multiple Routing Tables and Policy Routing
- Wireless updates
 - hardware crypto and fragmentation offload support
 - QoS (WME) support, "wireless spy support"
 - mixed PTK/GTK
 - CCMP/TKIP support and WE-19 HostAP support
 - BCM43xx wireless driver
 - ZD1211 wireless driver
 - WE-20, version 20 of the Wireless Extensions (2.6.17)
 - added the hardware-independent software MAC layer, "Soft MAC" (2.6.17)
 - added LEAP authentication type
- Added generic segmentation offload (GSO) (2.6.18)
 - can improve performance in some cases, though it needs to be enabled through `ethtool`
- DCCPv6 (2.6.16)

Added Hardware Support



NOTE

This section only enumerates the most generic features among many.

- x86-64 clustered APIC support (2.6.10)
- Infiniband support (2.6.11)
- Hot plug
 - added generic memory add/remove and supporting functions for memory hotplug (2.6.15)
- SATA/libata enhancements, additional hardware support
 - A completely reworked libata error handler; the result of all this work should be a more robust SATA subsystem which can recover from a wider range of errors.
 - Native Command Queuing (NCQ), the SATA version of tagged command queuing - the ability to have several I/O requests to the same drive outstanding at the same time. (2.6.18)
 - Hotplug support (2.6.18)
- EDAC support (2.6.16)

- The EDAC goal is to detect and report errors that occur within the system.
- Added a new ioatdma driver for the Intel(R) I/OAT DMA engine (2.6.18)

NUMA (Non-Uniform Memory Access) / Multi-core

- Cpusets (2.6.12)
 - Cpusets now provide a mechanism for assigning a set of CPUs and Memory Nodes to a set of tasks. Cpusets constrain the CPU and memory placement of tasks only to the resources within a task's current cpuset. These are essential in managing dynamic job placement on large systems.
- NUMA-aware slab allocator (2.6.14)
 - This creates slabs on multiple nodes and manages slabs in such a way that locality of allocations is optimized. Each node has its own list of partial, free and full slabs. All object allocations for a node occur from node-specific slab lists.
- Swap migration (2.6.16)
 - Swap migration allows the moving of physical location of pages between nodes in a NUMA system while the process is running.
- Huge pages (2.6.16)
 - Added NUMA policy support for huge pages: the `huge_zonelist()` function in the memory policy layer provides a list of zones ordered by NUMA distance. The `hugetlb` layer will walk that list looking for a zone that has available huge pages but is also in the `nodeset` of the current cpuset.
 - Huge pages now obey cpusets.
- Per-zone VM counters
 - provide zone-based VM statistics, which are necessary in determining what state of memory a zone is in
- Netfilter `ip_tables`: NUMA-aware allocation. (2.6.16)
- Multi-core
 - Added a new scheduler domain for representing multi-core with shared caches between cores. This makes it possible to make smarter cpu scheduling decisions on such systems, improving performance greatly for some cases. (2.6.17)
 - Power saving policy for the CPU scheduler: with `multicore/smt` cpus, the power consumption can be improved by leaving some packages idle while others do all the work, instead of spreading the tasks over all CPUs.