



# Red Hat Enterprise Linux 6

## Developer Guide

An introduction to application development tools in Red Hat Enterprise Linux 6



# Red Hat Enterprise Linux 6 Developer Guide

---

An introduction to application development tools in Red Hat Enterprise Linux 6

Robert Krátky  
Red Hat Customer Content Services  
rkratky@redhat.com

Don Domingo  
Red Hat Customer Content Services

Jacquelynn East  
Red Hat Customer Content Services

## Legal Notice

Copyright © 2016 Red Hat, Inc. and others.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This document describes the different features and utilities that make Red Hat Enterprise Linux 6 an ideal enterprise platform for application development.

## Table of Contents

|  |          |
|--|----------|
| <b>CHAPTER 1. COLLABORATING</b> .....        | <b>6</b> |
| 1.1. GIT                                     | 6        |
| 1.1.1. Installing and Configuring Git        | 6        |
| Installing the git Package                   | 6        |
| Configuring the Default Text Editor          | 6        |
| Setting Up User Information                  | 6        |
| 1.1.2. Creating a New Repository             | 7        |
| Initializing an Empty Repository             | 7        |
| Importing Data to a Repository               | 7        |
| 1.1.3. Cloning an Existing Repository        | 7        |
| 1.1.4. Adding, Renaming, and Deleting Files  | 8        |
| Adding Files and Directories                 | 8        |
| Renaming Files and Directories               | 8        |
| Deleting Files and Directories               | 8        |
| 1.1.5. Viewing Changes                       | 9        |
| Viewing the Current Status                   | 9        |
| Viewing Differences                          | 9        |
| 1.1.6. Committing Changes                    | 9        |
| 1.1.7. Sharing Changes                       | 9        |
| Pushing Changes to a Public Repository       | 10       |
| Creating Patches from Individual Commits     | 10       |
| 1.1.8. Updating a Repository                 | 10       |
| 1.1.9. Additional Resources                  | 10       |
| Installed Documentation                      | 10       |
| Online Documentation                         | 11       |
| 1.2. APACHE SUBVERSION (SVN)                 | 11       |
| 1.2.1. Installing and Configuring Subversion | 11       |
| Installing the subversion Package            | 11       |
| Setting Up the Default Editor                | 11       |
| 1.2.2. Creating a New Repository             | 11       |
| Initializing an Empty Repository             | 12       |
| Importing Data to a Repository               | 12       |
| 1.2.3. Checking Out a Working Copy           | 12       |
| 1.2.4. Adding, Renaming, and Deleting Files  | 13       |
| Adding a File or Directory                   | 13       |
| Renaming a File or Directory                 | 14       |
| Deleting a File or Directory                 | 14       |
| 1.2.5. Viewing Changes                       | 15       |
| Viewing the Status                           | 15       |
| Viewing Differences                          | 15       |
| 1.2.6. Committing Changes                    | 16       |
| 1.2.7. Updating a Working Copy               | 17       |
| 1.2.8. Additional Resources                  | 17       |
| Installed Documentation                      | 17       |
| Online Documentation                         | 18       |
| 1.3. CONCURRENT VERSIONS SYSTEM (CVS)        | 18       |
| 1.3.1. Installing and Configuring CVS        | 18       |
| Installing the cvs Package                   | 18       |
| Setting Up the Default Editor                | 18       |
| 1.3.2. Creating a New Repository             | 18       |
| Initializing an Empty Repository             | 19       |

|   |           |
|---|-----------|
| Importing Data to a Repository                        | 19        |
| 1.3.3. Checking Out a Working Copy                    | 20        |
| 1.3.4. Adding and Deleting Files                      | 20        |
| Adding a File   | 20        |
| Deleting a File                                       | 21        |
| 1.3.5. Viewing Changes                                | 21        |
| Viewing the Status                                    | 21        |
| Viewing Differences                                   | 22        |
| 1.3.6. Committing Changes                             | 23        |
| 1.3.7. Updating a Working Copy                        | 23        |
| 1.3.8. Additional Resources                           | 24        |
| Installed Documentation                               | 24        |
| <b>CHAPTER 2. LIBRARIES AND RUNTIME SUPPORT</b> ..... | <b>25</b> |
| 2.1. COMPATIBILITY                                    | 25        |
| 2.1.1. Static Linking                                 | 25        |
| 2.2. LIBRARY AND RUNTIME DETAILS                      | 26        |
| 2.2.1. The GNU C++ Standard Library                   | 26        |
| 2.2.1.1. Additional information                       | 26        |
| 2.2.2. Boost  | 26        |
| 2.2.2.1. Additional Information                       | 27        |
| 2.2.3. Qt   | 27        |
| 2.2.3.1. Qt Updates                                   | 27        |
| 2.2.3.2. Qt Creator                                   | 28        |
| 2.2.3.3. Qt Library Documentation                     | 28        |
| 2.2.4. KDE Development Framework                      | 28        |
| 2.2.4.1. KDE4 Architecture                            | 29        |
| 2.2.4.2. kdelibs Documentation                        | 30        |
| 2.2.5. GNOME Power Manager                            | 30        |
| 2.2.5.1. GNOME Power Management Version Guide         | 30        |
| 2.2.5.2. API Changes for glib                         | 31        |
| 2.2.5.3. API Changes for GTK+                         | 32        |
| 2.2.6. NSS Shared Databases                           | 34        |
| 2.2.6.1. Backwards Compatibility                      | 34        |
| 2.2.6.2. NSS Shared Databases Documentation           | 34        |
| 2.2.7. Python   | 34        |
| 2.2.7.1. Python Updates                               | 35        |
| 2.2.7.2. Python Documentation                         | 35        |
| 2.2.8. Java   | 36        |
| 2.2.8.1. Java Documentation                           | 36        |
| 2.2.9. Ruby   | 36        |
| 2.2.9.1. Ruby Documentation                           | 37        |
| 2.2.10. Perl  | 37        |
| 2.2.10.1. Perl Updates                                | 37        |
| 2.2.10.2. Installation                                | 39        |
| 2.2.10.3. Perl Documentation                          | 40        |
| <b>CHAPTER 3. COMPILING AND BUILDING</b> .....        | <b>42</b> |
| 3.1. GNU COMPILER COLLECTION (GCC)                    | 42        |
| 3.1.1. Language Compatibility                         | 42        |
| 3.1.2. Object Compatibility and Interoperability      | 44        |
| 3.1.3. Running GCC                                    | 45        |
| 3.1.3.1. Simple C Usage                               | 45        |

|   |           |
|---|-----------|
| 3.1.3.2. Simple C++ Usage                                       | 45        |
| 3.1.3.3. Simple Multi-File Usage                                | 46        |
| 3.1.3.4. Recommended Optimization Options                       | 47        |
| 3.1.3.5. Using Profile Feedback to Tune Optimization Heuristics | 48        |
| 3.1.3.6. Using 32-bit compilers on a 64-bit host                | 49        |
| 3.1.4. GCC Documentation  | 51        |
| 3.2. AUTOTOOLS  | 51        |
| 3.2.1. Autotools Plug-in for Eclipse                            | 52        |
| 3.2.2. Configuration Script                                     | 52        |
| 3.2.3. Autotools Documentation                                  | 52        |
| 3.3. BUILD-ID UNIQUE IDENTIFICATION OF BINARIES                 | 53        |
| <b>CHAPTER 4. DEBUGGING</b> .....                               | <b>54</b> |
| 4.1. ELF EXECUTABLE BINARIES                                    | 54        |
| 4.2. INSTALLING DEBUGINFO PACKAGES                              | 55        |
| 4.2.1. Installing Debuginfo Packages for Core Files Analysis    | 56        |
| 4.3. GDB  | 58        |
| 4.3.1. Simple GDB   | 59        |
| 4.3.2. Running GDB  | 61        |
| 4.3.3. Conditional Breakpoints                                  | 62        |
| 4.3.4. Forked Execution   | 63        |
| 4.3.5. Debugging Individual Threads                             | 65        |
| 4.3.6. Alternative User Interfaces for GDB                      | 69        |
| 4.4. VARIABLE TRACKING AT ASSIGNMENTS                           | 70        |
| 4.5. PYTHON PRETTY-PRINTERS                                     | 70        |
| <b>CHAPTER 5. PROFILING</b> .....                               | <b>74</b> |
| 5.1. VALGRIND   | 74        |
| 5.1.1. Valgrind Tools   | 74        |
| 5.1.2. Using Valgrind   | 75        |
| 5.1.3. Additional information                                   | 75        |
| 5.2. OPROFILE   | 75        |
| 5.2.1. Using OProfile   | 76        |
| 5.2.2. OProfile in Red Hat Enterprise Linux 7                   | 78        |
| 5.2.2.1. New Features   | 78        |
| 5.2.2.2. Known Problems and Limitations                         | 78        |
| 5.2.3. OProfile Documentation                                   | 79        |
| 5.3. SYSTEMTAP  | 79        |
| 5.3.1. Additional Information                                   | 79        |
| 5.4. PERFORMANCE COUNTERS FOR LINUX (PCL) TOOLS AND PERF        | 80        |
| 5.4.1. Perf Tool Commands                                       | 80        |
| 5.4.2. Using Perf   | 80        |
| 5.5. FTRACE   | 83        |
| 5.5.1. Using ftrace   | 83        |
| 5.5.2. ftrace Documentation                                     | 84        |
| <b>CHAPTER 6. DOCUMENTATION TOOLS</b> .....                     | <b>85</b> |
| 6.1. DOXYGEN  | 85        |
| 6.1.1. Doxygen Supported Output and Languages                   | 85        |
| 6.1.2. Getting Started  | 85        |
| 6.1.3. Running Doxygen  | 86        |
| 6.1.4. Documenting the Sources                                  | 87        |
| 6.1.5. Resources  | 91        |

|   |           |
|---|-----------|
| <b>APPENDIX A. APPENDIX</b> .....         | <b>92</b> |
| A.1. MALLOPT                              | 92        |
| malloc_trim                               | 92        |
| malloc_stats                              | 93        |
| Further Information                       | 93        |
| <b>APPENDIX B. REVISION HISTORY</b> ..... | <b>94</b> |
| <b>INDEX</b> .....                        | <b>95</b> |





## CHAPTER 1. COLLABORATING

Effective revision control is essential to all multi-developer projects. It allows all developers in a team to create, review, revise, and document code in a systematic and orderly manner. Red Hat Enterprise Linux 6 supports three of the most popular open-source revision control systems: **Git**, **SVN**, and **CVS**.

The following sections provide a brief overview and references to relevant documentation for each tool.

### 1.1. GIT

**Git** is a *distributed revision control system* with a peer-to-peer architecture. As opposed to centralized version control systems with a client-server model, **Git** ensures that each working copy of a **Git** repository is its exact copy with complete revision history. This not only allows you to work on and contribute to projects without the need to have permission to push your changes to their official repositories, but also makes it possible for you to work with no network connection.

#### 1.1.1. Installing and Configuring Git

##### Installing the git Package

In Red Hat Enterprise Linux 6, **Git** is provided by the `git` package. To install the `git` package and all its dependencies on your system, type the following at a shell prompt as **root**:

```
~]# yum install git
```

##### Configuring the Default Text Editor

Certain **Git** commands, such as `git commit`, require the user to write a short message or make some changes in an external text editor. To determine which text editor to start, **Git** attempts to read the value of the `GIT_EDITOR` environment variable, the `core.editor` configuration option, the `VISUAL` environment variable, and finally the `EDITOR` environment variable in this particular order. If none of these options and variables are specified, the `git` command starts `vi` as a reasonable default option.

To change the value of the `core.editor` configuration option in order to specify a different text editor, type the following at a shell prompt:

```
git config --global core.editor command
```

Replace *command* with the command to be used to start the selected text editor.

##### Example 1.1. Configuring the Default Text Editor

To configure **Git** to use `vim` as the default text editor, type the following at a shell prompt:

```
~]$ git config --global core.editor vim
```

##### Setting Up User Information

In **Git**, each commit (or revision) is associated with the full name and email of the person responsible for it. By default, **Git** uses an identity based on the user name and the host name.

To change the full name associated with your **Git** commits, type the following at a shell prompt:

```
git config --global user.name "full name"
```

To change the email address associated with your **Git** commits, type:

```
git config --global user.email "email_address"
```

### Example 1.2. Setting Up User Information

To configure **Git** to use **John Doe** as your full name and **john@example.com** as your email address, type the following at a shell prompt:

```
~]$ git config --global user.name "John Doe"
~]$ git config --global user.email "john@example.com"
```

## 1.1.2. Creating a New Repository

A *repository* is a place where **Git** stores all files that are under revision control, as well as additional data related to these files, such as the complete history of changes or information about who made those changes and when. Unlike in centralized revision control systems like Subversion or CVS, a **Git** repository and a *working directory* are usually the same. A typical **Git** repository also only stores a single project and when publicly accessible, it allows anyone to create its clone with a complete revision history.

### Initializing an Empty Repository

To create a new, empty **Git** repository, change to the directory in which you want to keep the repository and type the following at a shell prompt:

```
git init
```

This creates a hidden directory named **.git** in which all repository information is stored.

### Importing Data to a Repository

To put an existing project under revision control, create a **Git** repository in the directory with the project and run the following command:

```
git add .
```

This marks all files and directories in the current working directory as ready to be added to the **Git** repository. To proceed and actually add this content to the repository, commit the changes by typing the following at a shell prompt:

```
git commit [-m "commit message"]
```

Replace *commit message* with a short description of your revision. If you omit the **-m** option, this command allows you to write the commit message in an external text editor. For information on how to configure the default text editor, see [the section called “Configuring the Default Text Editor”](#).

## 1.1.3. Cloning an Existing Repository

To clone an existing **Git** repository, type the following at a shell prompt:

```
git clone git_repository [directory]
```

Replace *git\_repository* with a URL or a path to the **Git** repository you want to clone, and *directory* with a path to the directory in which you want to store the clone.

### 1.1.4. Adding, Renaming, and Deleting Files

#### Adding Files and Directories

To add an existing file to a **Git** repository and put it under revision control, change to the directory with your local **Git** repository and type the following at a shell prompt:

```
git add file...
```

Replace *file* with the file or files you want to add. This command marks the selected file or files as ready to be added to the **Git** repository. Similarly, to add all files that are stored in a certain directory to a **Git** repository, type:

```
git add directory...
```

Replace *directory* with the directory or directories you want to add. This command marks all files in the selected directory or directories as ready to be added to the **Git** repository.

To proceed and actually add this content to the repository, commit the changes as described in [Section 1.1.6, “Committing Changes”](#).

#### Renaming Files and Directories

To rename an existing file or directory in a **Git** repository, change to the directory with your local **Git** repository and type the following at a shell prompt:

```
git mv old_name new_name
```

Replace *old\_name* with the current name of the file or directory and *new\_name* with the new name. This command renames the selected file or directory and marks it as ready to be renamed in the **Git** repository.

To proceed and actually rename the content in the repository, commit the changes as described in [Section 1.1.6, “Committing Changes”](#).

#### Deleting Files and Directories

To delete an existing file from a **Git** repository, change to the directory with your local **Git** repository and type the following at a shell prompt:

```
git rm file...
```

Replace *file* with the file or files you want to delete. This command deletes all selected files and marks them as ready to be deleted from the **Git** repository. Similarly, to delete all files that are stored in a certain directory from a **Git** repository, type:

```
git rm -r directory...
```

Replace *directory* with the directory or directories you want to delete. This command deletes all selected directories and marks them as ready to be deleted from the **Git** repository.

To proceed and actually delete this content from the repository, commit the changes as described in [Section 1.1.6, “Committing Changes”](#).

## 1.1.5. Viewing Changes

### Viewing the Current Status

To determine the current status of your local **Git** repository, change to the directory with the repository and type the following command at a shell prompt:

```
git status
```

This command displays information about all uncommitted changes in the repository (**new file**, **renamed**, **deleted**, or **modified**) and tells you which changes will be applied the next time you commit them. For information on how to commit your changes, see [Section 1.1.6, “Committing Changes”](#).

### Viewing Differences

To view all changes in a **Git** repository, change to the directory with the repository and type the following at a shell prompt:

```
git diff
```

This command displays changes between the files in the repository and their latest revision. If you are only interested in changes in a particular file, supply its name on the command line as follows:

```
git diff file...
```

Replace *file* with the file or files you want to view.

## 1.1.6. Committing Changes

To apply your changes to a **Git** repository and create a new revision, change to the directory with the repository and type the following command at a shell prompt:

```
git commit [-m "commit message"]
```

Replace *commit message* with a short description of your revision. This command commits all changes in files that are explicitly marked as ready to be committed. To commit changes in all files that are under revision control, add the **-a** command line option as follows:

```
git commit -a [-m "commit message"]
```

Note that if you omit the **-m** option, the command allows you to write the commit message in an external text editor. For information on how to configure the default text editor, see [the section called “Configuring the Default Text Editor”](#).

## 1.1.7. Sharing Changes

Unlike in centralized version control systems such as CVS or Subversion, when working with **Git**, project contributors usually do not make their changes in a single, central repository. Instead, they either create a publicly accessible clone of their local repository, or submit their changes to others over email as patches.

## Pushing Changes to a Public Repository

To push your changes to a publicly accessible **Git** repository, change to the directory with your local repository and type the following at a shell prompt:

```
git push remote_repository
```

Replace *remote\_repository* with the name of the remote repository you want to push your changes to. Note that the repository from which you originally cloned your local copy is automatically named **origin**.

## Creating Patches from Individual Commits

To create patches from your commits, change to the directory with your local **Git** repository and type the following at a shell prompt:

```
git format-patch remote_repository
```

Replace *remote\_repository* with the name of the remote repository from which you made your local copy. This creates a patch for each commit that is not present in this remote repository.

### 1.1.8. Updating a Repository

To update your local copy of a **Git** repository and get the latest changes from a remote repository, change to the directory with your local **Git** repository and type the following at a shell prompt:

```
git fetch remote_repository
```

Replace *remote\_repository* with the name of the remote repository. This command fetches information about the current status of the remote repository, allowing you to review these changes before applying them to your local copy. To proceed and merge these changes with what you have in your local **Git** repository, type:

```
git merge remote_repository
```

Alternatively, you can perform both these steps at the same time by using the following command instead:

```
git pull remote_repository
```

### 1.1.9. Additional Resources

A detailed description of **Git** and its features is beyond the scope of this book. For more information about this revision control system, see the resources listed below.

#### Installed Documentation

- **gittutorial(7)** — The manual page named **gittutorial** provides a brief introduction to **Git** and its usage.
- **gittutorial-2(7)** — The manual page named **gittutorial-2** provides the second part of a brief introduction to **Git** and its usage.
- *Git User's Manual* — HTML documentation for **Git** is located at `/usr/share/doc/git-1.7.1/user-manual.html`.

## Online Documentation

- [Pro Git](#) — The online version of the *Pro Git* book provides a detailed description of **Git**, its concepts and its usage.

## 1.2. APACHE SUBVERSION (SVN)

**Apache Subversion**, commonly abbreviated as **SVN**, is a *centralized version control system* with a client-server architecture. It is a successor to the older Concurrent Versions System (CVS), preserves the same development model, and addresses problems often encountered with CVS.

### 1.2.1. Installing and Configuring Subversion

#### Installing the subversion Package

In Red Hat Enterprise Linux 6, **Subversion** is provided by the `subversion` package. To install the subversion package and all its dependencies on your system, type the following at a shell prompt as **root**:

```
yum install subversion
```

This installs a command line Subversion client, a Subversion server, and other related tools to the system.

#### Setting Up the Default Editor

When using Subversion on the command line, certain commands such as `svn import` or `svn commit` require the user to write a short log message. To determine which text editor to start, the `svn` client application first reads the contents of the environment variable `$SVN_EDITOR`, then reads more general environment variables `$VISUAL` and `$EDITOR`, and if none of these is set, it reports an error.

To persistently change the value of the `$SVN_EDITOR` environment variable, run the following command:

```
echo "export SVN_EDITOR=command" >> ~/.bashrc
```

This adds the `export SVN_EDITOR=command` line to your `~/.bashrc` file. Replace `command` with a command that runs the editor of your choice (for example, `emacs`). Note that for this change to take effect in the current shell session, you must execute the commands in `~/.bashrc` by typing the following at a shell prompt:

```
. ~/.bashrc
```

#### Example 1.3. Setting up the default text editor

To configure the Subversion client to use **Emacs** as a text editor, type:

```
~]$ echo "export SVN_EDITOR=emacs" >> ~/.bashrc
~]$ . ~/.bashrc
```

### 1.2.2. Creating a New Repository

A Subversion *repository* is a central place to store files and directories that are under revision control, as well as additional data such as a complete history of changes or information about who made those

changes and when. A typical Subversion repository stores multiple projects in separate subdirectories. When publicly accessible, it allows several developers to create a *working copy* of any of the subdirectories, make changes, and share these changes with others by *committing* them back to the repository.

### Initializing an Empty Repository

To create a new, empty Subversion repository in a directory of your choice, run the following command:

```
svnadmin create path
```

Note that *path* is an absolute or relative path to the directory in which you want to store the repository (for example, `/var/svn/`). If the directory does not exist, `svnadmin create` creates it for you.

#### Example 1.4. Initializing a new Subversion repository

To create an empty Subversion repository in the `~/svn/` directory, type:

```
~]$ svnadmin create svn
```

### Importing Data to a Repository

To put an existing project under revision control, run the following command:

```
svn import local_path svn_repository/remote_path [-m "commit message"]
```

Note that *local\_path* is an absolute or relative path to the directory in which you keep the project (use `.` for the current working directory), *svn\_repository* is a URL of the Subversion repository, and *remote\_path* is the target directory in the Subversion repository (for example, `project/trunk`).

#### Example 1.5. Importing a project to a Subversion repository

Imagine that the directory with your project has the following contents:

```
~]$ ls myproject
AUTHORS  doc  INSTALL  LICENSE  Makefile  README  src  TODO
```

Also imagine that you have an empty Subversion repository in `~/svn/` (in this example, `/home/john/svn/`). To import the project under `project/trunk` in this repository, type:

```
~]$ svn import myproject file:///home/john/svn/project/trunk -m "Initial
import."
Adding      project/AUTHORS
Adding      project/doc
Adding      project/doc/index.html
Adding      project/INSTALL
Adding      project/src
...
```

### 1.2.3. Checking Out a Working Copy

To check out a working copy of a project in a Subversion repository, run the following command:



```
svn checkout svn_repository/remote_path [directory]
```

This creates a new directory called *directory* with a working copy of a project in it. Note that *svn\_repository* is a URL of the Subversion repository and *remote\_path* is the subdirectory in which the project is stored.

### Example 1.6. Checking out a working copy

Imagine that you have a Subversion repository in the `~/svn/` directory (in this case, `/home/john/svn/`) and that this repository contains the latest version of a project in the `project/trunk` subdirectory. To check out a working copy of this project, type:

```
~]$ svn checkout svn:///home/john/svn/project/trunk project
A   project/AUTHORS
A   project/doc
A   project/doc/index.html
A   project/INSTALL
A   project/src
...
```

## 1.2.4. Adding, Renaming, and Deleting Files

### Adding a File or Directory

To add an existing file to a Subversion repository and put it under revision control, change to the directory with its working copy and run the following command:

```
svn add file...
```

Similarly, to add a directory and all files that are in it, type:

```
svn add directory...
```

This schedules the files and directories for addition to the Subversion repository. To proceed and actually add this content to the repository, run the `svn commit` command as described in [Section 1.2.6](#), “Committing Changes”.

### Example 1.7. Adding a file to a Subversion repository

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  doc  INSTALL  LICENSE  Makefile  README  src  TODO
```

With the exception of `ChangeLog`, all files and directories within this directory are already under revision control. To schedule this file for addition to the Subversion repository, type:

```
project]$ svn add ChangeLog
A   ChangeLog
```

## Renaming a File or Directory

To rename an existing file or directory in a Subversion repository, change to the directory with its working copy and run the following command:

```
svn move old_name new_name
```

This creates a duplicate of the original file or directory, schedules it for addition, and automatically deletes the original. To proceed and actually rename the content in the Subversion repository, run the `svn commit` command as described in [Section 1.2.6, “Committing Changes”](#).

### Example 1.8. Renaming a file in a Subversion repository

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  doc  INSTALL  LICENSE  Makefile  README  src  TODO
```

All files in this directory are under revision control. To schedule the **LICENSE** file for renaming to **COPYING**, type:

```
project]$ svn move LICENSE COPYING
A      COPYING
D      LICENSE
```

Note that `svn move` automatically renames the file in your working copy:

```
project]$ ls
AUTHORS  ChangeLog  COPYING  doc  INSTALL  Makefile  README  src  TODO
```

## Deleting a File or Directory

To remove a file from a Subversion repository, change to the directory with its working copy and run the following command:

```
svn delete file...
```

Similarly, to remove a directory and all files that are in it, type:

```
svn delete directory...
```

This schedules the files and directories for removal from the Subversion repository. To proceed and actually remove this content from the repository, run the `svn commit` command as described in [Section 1.2.6, “Committing Changes”](#).

### Example 1.9. Deleting a file from a Subversion repository

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  COPYING  doc  INSTALL  Makefile  README  src  TODO
```

All files in this directory are under revision control. To schedule the **TODO** file for removal from the SVN repository, type:

```
project]$ svn delete TODO
D      TODO
```

Note that **svn delete** automatically deletes the file from your working copy:

```
project]$ ls
AUTHORS  ChangeLog  COPYING  doc  INSTALL  Makefile  README  src
```

## 1.2.5. Viewing Changes

### Viewing the Status

To determine the current status of a working copy, change to the directory with the working copy and run the following command:

```
svn status
```

This displays information about all changes to the working copy (**A** for a file that is scheduled for addition, **D** for a file that is scheduled for removal, **M** for a file that contains local changes, **C** for a file with unresolved conflicts, **?** for a file that is not under revision control).

#### Example 1.10. Viewing the status of a working copy

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  COPYING  doc  INSTALL  Makefile  README  src
```

With the exception of **ChangeLog**, which is scheduled for addition to the Subversion repository, all files and directories within this directory are already under revision control. The **TODO** file, which is also under revision control, has been scheduled for removal and is no longer present in the working copy. The **LICENSE** file has been renamed to **COPYING**, and **Makefile** contains local changes. To display the status of such a working copy, type:

```
project]$ svn status
D      LICENSE
D      TODO
A      ChangeLog
A +    COPYING
M      Makefile
```

### Viewing Differences

To view differences between a working copy and the checked out content, change to the directory with the working copy and run the following command:

```
svn diff [file...]
```

This displays changes to all files in the working copy. If you are only interested in changes to a particular file, supply the file name on the command line.

### Example 1.11. Viewing changes to a working copy

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  COPYING  CVS  doc  INSTALL  Makefile  README  src
```

All files in this directory are under revision control and **Makefile** contains local changes. To view these changes, type:

```
project]$ svn diff Makefile
Index: Makefile
=====
--- Makefile      (revision 1)
+++ Makefile      (working copy)
@@ -153,7 +153,7 @@
     -rmdir $(man1dir)

clean:
-       -rm -f $(MAN1)
+       -rm -f $(MAN1) $(MAN7)

%.1: %.pl
        $(POD2MAN) --section=1 --release="Version $(VERSION)" \
```

## 1.2.6. Committing Changes

To share your changes with others and commit them to a Subversion repository, change to the directory with its working copy and run the following command:

```
svn commit [-m "commit message"]
```

Note that unless you specify the commit message on the command line, Subversion opens an external text editor for you to write it. For information on how to determine which editor to start, see [Section 1.2.1, "Installing and Configuring Subversion"](#).

### Example 1.12. Committing changes to a Subversion repository

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  COPYING  doc  INSTALL  Makefile  README  src
```

In this working copy, **ChangeLog** is scheduled for addition to the Subversion repository, **Makefile** already is under revision control and contains local changes, and the **TODO** file, which is also under revision control, has been scheduled for removal and is no longer present in the working copy.

Additionally, the **LICENSE** file has been renamed to **COPYING**. To commit these changes to the Subversion repository, type:

```
project]$ svn commit -m "Updated the makefile."
Adding          COPYING
Adding          ChangeLog
Deleting        LICENSE
Sending         Makefile
Deleting        TODO
Transmitting file data ..
Committed revision 2.
```

## 1.2.7. Updating a Working Copy

To update a working copy and get the latest changes from a Subversion repository, change to the directory with the working copy and run the following command:

```
svn update
```

### Example 1.13. Updating a working copy

Imagine that the directory with your working copy of a Subversion repository has the following contents:

```
project]$ ls
AUTHORS  doc  INSTALL  LICENSE  Makefile  README  src  TODO
```

Also imagine that somebody recently added **ChangeLog** to the repository, removed the **TODO** file from it, changed the name of **LICENSE** to **COPYING**, and made some changes to **Makefile**. To update this working copy, type:

```
myproject]$ svn update
D    LICENSE
D    TODO
A    COPYING
A    Changelog
M    Makefile
Updated to revision 2.
```

## 1.2.8. Additional Resources

A detailed description of all supported features is beyond the scope of this book. For more information, see the resources listed below.

### Installed Documentation

- **svn help** — The output of the **svn help** command provides detailed information on the **svn** usage.
- **svnadmin help** — The output of the **svnadmin help** command provides detailed information on the **svnadmin** usage.

## Online Documentation

- [Version Control with Subversion](#) — The official Subversion website refers to the *Version Control with Subversion* manual, which provides an in-depth description of Subversion, its administration and its usage.

## 1.3. CONCURRENT VERSIONS SYSTEM (CVS)

**Concurrent Versions System**, commonly abbreviated as **CVS**, is a *centralized version control system* with a client-server architecture. It is a successor to the older Revision Control System (RCS), and allows multiple developers to cooperate on the same project while keeping track of every change made to the files that are under revision control.

### 1.3.1. Installing and Configuring CVS

#### Installing the cvs Package

In Red Hat Enterprise Linux 6, **CVS** is provided by the `cvs` package. To install the `cvs` package and all its dependencies on your system, type the following at a shell prompt as **root**:

```
yum install cvs
```

This installs a command line CVS client, a CVS server, and other related tools to the system.

#### Setting Up the Default Editor

When using CVS on the command line, certain commands such as `cvs import` or `cvs commit` require the user to write a short log message. To determine which text editor to start, the `cvs` client application first reads the contents of the environment variable `$CVSEEDITOR`, then reads the more general environment variable `$EDITOR`, and if none of these is set, it starts `vi`.

To persistently change the value of the `$CVSEEDITOR` environment variable, run the following command:

```
echo "export CVSEEDITOR=command" >> ~/.bashrc
```

This adds the `export CVSEEDITOR=command` line to your `~/.bashrc` file. Replace `command` with a command that runs the editor of your choice (for example, `emacs`). Note that for this change to take effect in the current shell session, you must execute the commands in `~/.bashrc` by typing the following at a shell prompt:

```
. ~/.bashrc
```

#### Example 1.14. Setting up the default text editor

To configure the CVS client to use **Emacs** as a text editor, type:

```
~]$ echo "export CVSEEDITOR=emacs" >> ~/.bashrc
~]$ . ~/.bashrc
```

### 1.3.2. Creating a New Repository

A CVS *repository* is a central place to store files and directories that are under revision control, as well as additional data such as a complete history of changes or information about who made those changes

and when. A typical CVS repository stores multiple projects in separate subdirectories called *modules*. When publicly accessible, it allows several developers to create a *working copy* of any of the modules, make changes, and share these changes with others by *committing* them back to the repository.

### Initializing an Empty Repository

To create a new, empty CVS repository in a directory of your choice, run the following command:

```
cvs -d path init
```

Note that *path* must be an absolute path to the directory in which you want to store the repository (for example, `/var/cvs/`). Alternatively, you can specify this path by changing the value of the `$CVSRROOT` environment variable:

```
export CVSRROOT=path
```

This allows you to omit the path from `cvs init` and other CVS-related commands:

```
cvs init
```

#### Example 1.15. Initializing a new CVS repository

To create an empty CVS repository in the `~/cvs/` directory, type:

```
~]$ export CVSRROOT=~/cvs
~]$ cvs init
```

### Importing Data to a Repository

To put an existing project under revision control, change to the directory in which the project is stored and run the following command:

```
cvs [-d cvs_repository] import [-m "commit message"] module vendor_tag  
release_tag
```

Note that *cvs\_repository* is a path to the CVS repository, *module* is the subdirectory in which you want to import the project (such as `project`), and *vendor\_tag* and *release\_tag* are vendor and release tags.

#### Example 1.16. Importing a project to a CVS repository

Imagine that the directory with your project has the following contents:

```
~]$ ls myproject
AUTHORS  doc  INSTALL  LICENSE  Makefile  README  src  TODO
```

Also imagine that you have an empty CVS repository in `~/cvs/`. To import the project under `project` in this repository with vendor tag `mycompany` and release tag `init`, type:

```
myproject]$ export CVSRROOT=~/cvs
myproject]$ cvs import -m "Initial import." project mycompany init
N project/Makefile
N project/AUTHORS
N project/LICENSE
```

```
N project/TODO
N project/INSTALL
...
```

### 1.3.3. Checking Out a Working Copy

To check out a working copy of a project in a CVS repository, run the following command:

```
cvs -d cvs_repository checkout module
```

This creates a new directory called *module* with a working copy of a project in it. Note that *cvs\_repository* is a URL of the CVS repository and *module* is the subdirectory in which the project is stored (such as **project**). Alternatively, you can set the **\$CVSRROOT** environment variable as follows:

```
export CVSRROOT=cvs_repository
```

Then you can use the **cvs checkout** command without the **-d** option:

```
cvs checkout module
```

#### Example 1.17. Checking out a working copy

Imagine that you have a CVS repository in `~/cvs/` and that this repository contains a module named **project**. To check out a working copy of this module, type:

```
~]$ export CVSRROOT=~/cvs
~]$ cvs checkout project
cvs checkout: Updating project
U project/AUTHORS
U project/INSTALL
U project/LICENSE
U project/Makefile
U project/TODO
```

### 1.3.4. Adding and Deleting Files

#### Adding a File

To add an existing file to a CVS repository and put it under revision control, change to the directory with its working copy and run the following command:

```
cvs add file..
```

This schedules the file for addition to the CVS repository. To proceed and actually add the file to the repository, run the **cvs commit** command as described in [Section 1.3.6, “Committing Changes”](#).

#### Example 1.18. Adding a file to a CVS repository

Imagine that the directory with your working copy of a CVS repository has the following contents:



```
project]$ ls
AUTHORS  ChangeLog  CVS  doc  INSTALL  LICENSE  Makefile  README  src
TODO
```

With the exception of **ChangeLog**, all files and directories within this directory are already under revision control. To schedule this file for addition to the CVS repository, type:

```
project]$ cvs add ChangeLog
cvs add: scheduling file `ChangeLog' for addition
cvs add: use 'cvs commit' to add this file permanently
```

### Deleting a File

To remove a file from a CVS repository, change to the directory with its working copy and delete it locally:

```
rm file...
```

Then schedule this file for removal by using the following command:

```
cvs remove file...
```

To proceed and actually remove the file from the repository, run the **cvs commit** command as described in [Section 1.3.6, “Committing Changes”](#).

#### Example 1.19. Removing a file from a CVS repository

Imagine that the directory with your working copy of a CVS repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  CVS  doc  INSTALL  LICENSE  Makefile  README  src
TODO
```

All files in this directory are under revision control. To schedule the **TODO** file for removal from the CVS repository, type:

```
project]$ rm TODO
project]$ cvs remove TODO
cvs remove: scheduling `TODO' for removal
cvs remove: use 'cvs commit' to remove this file permanently
```

## 1.3.5. Viewing Changes

### Viewing the Status

To determine the current status of a working copy, change to the directory with the working copy and run the following command:

```
cvs status
```

This displays detailed information about each file that is under revision control, including its current

status (such as **Up-to-date**, **Locally Added**, **Locally Removed**, or **Locally Modified**) and revision. However, if you are only interested in what has changed in your working copy, you can simplify the output by typing the following at a shell prompt:

```
cvs status 2>/dev/null | grep Status: | grep -v Up-to-date
```

### Example 1.20. Viewing the status of a working copy

Imagine that the directory with your working copy of a CVS repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  CVS  doc  INSTALL  LICENSE  Makefile  README  src
```

With the exception of **ChangeLog**, which is scheduled for addition to the CVS repository, all files and directories within this directory are already under revision control. The **TODO** file, which is also under revision control, has been scheduled for removal and is no longer present in the working copy. Finally, **Makefile** contains local changes. To display the status of such a working copy, type:

```
project]$ cvs status 2>/dev/null | grep Status: | grep -v Up-to-date
File: ChangeLog          Status: Locally Added
File: Makefile           Status: Locally Modified
File: no file TODO       Status: Locally Removed
```

## Viewing Differences

To view differences between a working copy and the checked out content, change to the directory with the working copy and run the following command:

```
cvs diff [file...]
```

This displays changes to all files in the working copy. If you are only interested in changes to a particular file, supply the file name on the command line.

### Example 1.21. Viewing changes to a working copy

Imagine that the directory with your working copy of a CVS repository has the following contents:

```
project]$ ls
AUTHORS  ChangeLog  CVS  doc  INSTALL  LICENSE  Makefile  README  src
```

All files in this directory are under revision control and **Makefile** contains local changes. To view these changes, type:

```
project]$ cvs diff
cvs diff: Diffing .
cvs diff: ChangeLog is a new entry, no comparison available
Index: Makefile
=====
RCS file: /home/john/cvs/project/Makefile,v
retrieving revision 1.1.1.1
diff -r1.1.1.1 Makefile
156c156
<      -rm -f $(MAN1)
```

```

---
>      -rm -f $(MAN1) $(MAN7)
cvs diff: TODO was removed, no comparison available
cvs diff: Diffing doc
...

```

### 1.3.6. Committing Changes

To share your changes with others and commit them to a CVS repository, change to the directory with its working copy and run the following command:

```
cvs commit [-m "commit message"]
```

Note that unless you specify the commit message on the command line, CVS opens an external text editor (**vi** by default) for you to write it. For information on how to determine which editor to start, see [Section 1.3.1, “Installing and Configuring CVS”](#).

#### Example 1.22. Committing changes to a CVS repository

Imagine that the directory with your working copy of a CVS repository has the following contents:

```

project]$ ls
AUTHORS  ChangeLog  CVS  doc  INSTALL  LICENSE  Makefile  README  src

```

In this working copy, **ChangeLog** is scheduled for addition to the CVS repository, **Makefile** already is under revision control and contains local changes, and the **TODO** file, which is also under revision control, has been scheduled for removal and is no longer present in the working copy. To commit these changes to the CVS repository, type:

```

project]$ cvs commit -m "Updated the makefile."
cvs commit: Examining .
cvs commit: Examining doc
...
RCS file: /home/john/cvsroot/project/ChangeLog,v
done
Checking in ChangeLog;
/home/john/cvsroot/project/ChangeLog,v <-- ChangeLog
initial revision: 1.1
done
Checking in Makefile;
/home/john/cvsroot/project/Makefile,v <-- Makefile
new revision: 1.2; previous revision: 1.1
done
Removing TODO;
/home/john/cvsroot/project/TOD0,v <-- TODO
new revision: delete; previous revision: 1.1.1.1
done

```

### 1.3.7. Updating a Working Copy

To update a working copy and get the latest changes from a CVS repository, change to the directory with the working copy and run the following command:

## **cvs update**

### **Example 1.23. Updating a working copy**

Imagine that the directory with your working copy of a CVS repository has the following contents:

```
project]$ ls
AUTHORS  CVS  doc  INSTALL  LICENSE  Makefile  README  src  TODO
```

Also imagine that somebody recently added **ChangeLog** to the repository, removed the **TODO** file from it, and made some changes to **Makefile**. To update this working copy, type:

```
myproject]$ cvs update
cvs update: Updating .
U ChangeLog
U Makefile
cvs update: TODO is no longer in the repository
cvs update: Updating doc
cvs update: Updating src
```

### **1.3.8. Additional Resources**

A detailed description of all supported features is beyond the scope of this book. For more information, see the resources listed below.

#### **Installed Documentation**

- `cvs(1)` — The manual page for the **cvs** client program provides detailed information on its usage.

## CHAPTER 2. LIBRARIES AND RUNTIME SUPPORT

Red Hat Enterprise Linux 6 supports the development of custom applications in a wide variety of programming languages using proven, industrial-strength tools. This chapter describes the runtime support libraries provided in Red Hat Enterprise Linux 6.

### 2.1. COMPATIBILITY

Compatibility specifies the portability of binary objects and source code across different instances of a computer operating environment. Officially, Red Hat supports current release and two consecutive prior versions. This means that applications built on Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5 will continue to run on Red Hat Enterprise Linux 6 as long as they comply with Red Hat guidelines (using the symbols that have been white-listed, for example).

Red Hat understands that as an enterprise platform, customers rely on long-term deployment of their applications. For this reason, applications built against C/C++ libraries with the help of compatibility libraries continue to be supported for ten years.

There are two types of compatibility:

#### Source Compatibility

Source compatibility specifies that code will compile and execute in a consistent and predictable way across different instances of the operating environment. This type of compatibility is defined by conformance with specified Application Programming Interfaces (APIs).

#### Binary Compatibility

Binary Compatibility specifies that compiled binaries in the form of executables and *Dynamic Shared Objects* (DSOs) will run correctly across different instances of the operating environment. This type of compatibility is defined by conformance with specified Application Binary Interfaces (ABIs).



#### NOTE

Compilation of code using `-devel` packages from the *Optional* channel is not supported by Red Hat because libraries in the *Optional* channel are only provided for internal use by the operating system. They are not intended for general application development.

Among other things, this means that security issues in these libraries are only addressed if base operating system components provided by Red Hat are affected. It also means that you may not be able to easily determine whether the built application only uses libraries supported for use by non-Red Hat applications. Therefore, it is recommended to build software on systems that do not have packages from *Optional* channels installed.

For further information regarding this and all levels of compatibility between core and non-core libraries, see [Red Hat Enterprise Linux Life Cycle](#) and the general [Red Hat Enterprise Linux Application Compatibility Policies](#).

#### 2.1.1. Static Linking

Static linking is emphatically discouraged for all Red Hat Enterprise Linux releases. Static linking causes far more problems than it solves, and should be avoided at all costs.

The main drawback of static linking is that it is only guaranteed to work on the system on which it was

built, and even then only until the next release of glibc or libstdc++ (in the case of C++). There is no forward or backward compatibility with a static build. Furthermore, any security fixes (or general-purpose fixes) in subsequent updates to the libraries will not be available unless the affected statically linked executables are re-linked.

A few more reasons why static linking should be avoided are:

- Larger memory footprint.
- Slower application startup time.
- Reduced glibc features with static linking.
- Security measures like load address randomization cannot be used.
- Dynamic loading of shared objects outside of glibc is not supported.



#### NOTE

The `compat-glibc` package is included with Red Hat Enterprise Linux 6, but it is not a runtime package and therefore not required for running anything. It is solely a development package, containing header files and dummy libraries for linking. This allows compiling and linking packages to run in older Red Hat Enterprise Linux versions (using `compat-gcc-*` against those headers and libraries). Running `rpm -qpi compat-glibc-*` will provide some information on how to use this package.

## 2.2. LIBRARY AND RUNTIME DETAILS

### 2.2.1. The GNU C++ Standard Library

The `libstdc++` package contains the **GNU C++ Standard Library**, which is an ongoing project to implement the **ISO 14882 Standard C++** library.

Installing the `libstdc++` package will provide just enough to satisfy link dependencies (that is, only shared library files). To make full use of all available libraries and header files for C++ development, you must install `libstdc++-devel` as well. The `libstdc++-devel` package also contains a GNU-specific implementation of the Standard Template Library (STL).

For Red Hat Enterprise Linux 4, 5, and 6, the C++ language and runtime implementation has remained stable, and thus no compatibility libraries are required for **libstdc++**.

#### 2.2.1.1. Additional information

To use the **man** pages for library components, install the `libstdc++-docs` package. This will provide **man** page information for nearly all resources provided by the library; for example, to view information about the **vector** container, use its fully-qualified component name: **man std::vector**.

The `libstdc++-docs` package also provides manuals and reference information in HTML form in the following directory: `/usr/share/doc/libstdc++-docs-version/html/spine.html`.

### 2.2.2. Boost

The `boost` package contains a large number of free peer-reviewed portable C++ source libraries. These libraries are suitable for tasks such as portable file-system access and time or date abstraction,

serialization, unit testing, thread creation and multi-process synchronization, parsing, graphing, regular expression manipulation, and many others.

Installing the boost package will provide just enough libraries to satisfy link dependencies (that is, only shared library files). To make full use of all available libraries and header files for C++ development, you must install boost-devel as well.

The **boost** package is actually a meta-package, containing many library sub-packages. These sub-packages can also be installed individually to provide finer inter-package dependency tracking.

The meta-package does not include dependencies for packages for static linking or packages that depend on the underlying Message Passing Interface (MPI) support.

MPI support is provided in two forms: one for the default Open MPI implementation (package boost-openmpi) and another for the alternate MPICH2 implementation (package boost-mpich2). The selection of the underlying MPI library in use is up to the user and depends on specific hardware details and user preferences. Please note that these packages can be installed in parallel because installed files have unique directory locations.

If static linkage cannot be avoided, the boost-static package will install the necessary static libraries. Both thread-enabled and single-threaded libraries are provided.

### 2.2.2.1. Additional Information

The boost-doc package provides manuals and reference information in HTML form located in the following directory: `/usr/share/doc/boost-doc-version/index.html`.

The main site for the development of Boost is <http://boost.org>.

### 2.2.3. Qt

The **qt** package provides the Qt (pronounced "cute") cross-platform application development framework used in the development of GUI programs. Aside from being a popular "widget toolkit", Qt is also used for developing non-GUI programs such as console tools and servers. Qt was used in the development of notable projects such as Google Earth, KDE, Opera, OPIE, VoxOx, Skype, VLC media player and VirtualBox. It is produced by Nokia's Qt Development Frameworks division, which came into being after Nokia's acquisition of the Norwegian company Trolltech, the original producer of Qt, on June 17, 2008.

Qt uses standard C++ but makes extensive use of a special pre-processor called the *Meta Object Compiler* (MOC) to enrich the language. Qt can also be used in other programming languages via language bindings. It runs on all major platforms and has extensive internationalization support. Non-GUI Qt features include SQL database access, XML parsing, thread management, network support, and a unified cross-platform API for file handling.

Distributed under the terms of the GNU Lesser General Public License (among others), Qt is free and open source software. The Red Hat Enterprise Linux 6 version of Qt supports a wide range of compilers, including the GCC C++ compiler and the Visual Studio suite.

#### 2.2.3.1. Qt Updates

Some of the improvements the Red Hat Enterprise Linux 6 version of Qt include:

- Advanced user experience
  - **Advanced Graphics Effects:** options for opacity, drop-shadows, blur, colorization, and other similar effects

- **Animation and State Machine:** create simple or complex animations without the hassle of managing complex code
- Gesture and multi-touch support
- Support for new platforms
  - Windows 7, Mac OSX 10.6, and other desktop platforms are now supported
  - Added support for mobile development; Qt is optimized for the upcoming Maemo 6 platform, and will soon be ported to Maemo 5. In addition, Qt now supports the Symbian platform, with integration for the S60 framework.
  - Added support for Real-Time Operating Systems such as QNX and VxWorks
- Improved performance, featuring added support for hardware-accelerated rendering (along with other rendering updates)
- Updated cross-platform IDE

For more details on updates to Qt included in Red Hat Enterprise Linux 6, see the following links:

- <http://doc.qt.nokia.com/4.6/qt4-6-intro.html>
- <http://doc.qt.nokia.com/4.6/qt4-intro.html>

### 2.2.3.2. Qt Creator

**Qt Creator** is a cross-platform IDE tailored to the requirements of Qt developers. It includes the following graphical tools:

- An advanced C++ code editor
- Integrated GUI layout and forms designer
- Project and build management tools
- Integrated, context-sensitive help system
- Visual debugger
- Rapid code navigation tools

### 2.2.3.3. Qt Library Documentation

The **qt-doc** package provides HTML manuals and references located in `/usr/share/doc/qt4/html/`. This package also provides the *Qt Reference Documentation*, which is an excellent starting point for development within the Qt framework.

You can also install further demos and examples from **qt-demos** and **qt-examples**. To get an overview of the capabilities of the Qt framework, see `/usr/bin/qtdemo-qt4` (provided by **qt-demos**).

### 2.2.4. KDE Development Framework



The **kde1ibs-devel** package provides the KDE libraries, which build on Qt to provide a framework for making application development easier. The KDE development framework also helps provide consistency across the KDE desktop environment.

### 2.2.4.1. KDE4 Architecture

The KDE development framework's architecture in Red Hat Enterprise Linux 6 uses KDE4, which is built on the following technologies:

#### Plasma

**Plasma** replaces KDesktop in KDE4. Its implementation is based on the **Qt Graphics View Framework**, which was introduced in Qt 4.2. For more information about **Plasma**, see <http://techbase.kde.org/Development/Architecture/KDE4/Plasma>.

#### Sonnet

**Sonnet** is a multilingual spell-checking application that supports automatic language detection, primary/backup dictionaries, and other useful features. It replaces **kspe112** in KDE4.

#### KIO

The KIO library provides a framework for network-transparent file handling, allowing users to easily access files through network-transparent protocols. It also helps provides standard file dialogs.

#### KJS/KHTML

KJS and KHTML are fully-fledged JavaScript and HTML engines used by different applications native to KDE4 (such as **konqueror**).

#### Solid

**Solid** is a hardware and network awareness framework that allows you to develop applications with hardware interaction features. Its comprehensive API provides the necessary abstraction to support cross-platform application development. For more information, see <http://techbase.kde.org/Development/Architecture/KDE4/Solid>.

#### Phonon

**Phonon** is a multimedia framework that helps you develop applications with multimedia functionalities. It facilitates the usage of media capabilities within KDE. For more information, see <http://techbase.kde.org/Development/Architecture/KDE4/Phonon>.

#### Telepathy

**Telepathy** provides a real-time communication and collaboration framework within KDE4. Its primary function is to tighten integration between different components within KDE. For a brief overview on the project, see [http://community.kde.org/Real-Time\\_Communication\\_and\\_Collaboration](http://community.kde.org/Real-Time_Communication_and_Collaboration).

#### Akonadi

**Akonadi** provides a framework for centralizing storage of *Parallel Infrastructure Management* (PIM) components. For more information, see <http://techbase.kde.org/Development/Architecture/KDE4/Akonadi>.

#### Online Help within KDE4

KDE4 also features an easy-to-use Qt-based framework for adding online help capabilities to applications. Such capabilities include tooltips, hover-help information, and **khelphcenter** manuals.

For a brief overview on online help within KDE4, see [http://techbase.kde.org/Development/Architecture/KDE4/Providing\\_Online\\_Help](http://techbase.kde.org/Development/Architecture/KDE4/Providing_Online_Help).

## KXMLGUI

**KXMLGUI** is a framework for designing user interfaces using XML. This framework allows you to design UI elements based on "actions" (defined by the developer) without having to revise source code. For more information, see [https://techbase.kde.org/Development/Architecture/KDE3/XMLGUI\\_Technology](https://techbase.kde.org/Development/Architecture/KDE3/XMLGUI_Technology).

## Strigi

**Strigi** is a desktop search daemon compatible with many desktop environments and operating systems. It uses its own **jstream** system which allows for deep indexing of files. For more information on the development of **Strigi**, see <http://www.vandenoever.info/software/strigi/>.

## KNewStuff2

**KNewStuff2** is a collaborative data sharing library used by many KDE4 applications. For more information, see <http://techbase.kde.org/Projects/KNS2>.

### 2.2.4.2. kdelibs Documentation

The **kdelibs-apidocs** package provides HTML documentation for the KDE development framework in `/usr/share/doc/HTML/en/kdelibs4-apidocs/`. The following links also provide details on KDE-related programming tasks:

- <http://techbase.kde.org/>
- <http://techbase.kde.org/Development/Tutorials>
- <http://techbase.kde.org/Development/FAQs>
- <http://api.kde.org>

### 2.2.5. GNOME Power Manager

The backend program of the GNOME power management infrastructure is **gnome-power-manager**. It was introduced in Red Hat Enterprise Linux 5 and provides a complete and integrated solution to power management under the GNOME desktop environment. In Red Hat Enterprise Linux 6, the storage-handling parts of **hal** was replaced by **udisks**, and the **libgnomeprint** stack was replaced by print support in **gtk2**.

#### 2.2.5.1. GNOME Power Management Version Guide

This section will detail what versions of **gnome-power-management** are shipped with the various Red Hat Enterprise Linux versions.

In general, however, Red Hat Enterprise Linux 4 ships with GNOME 2.8, Red Hat Enterprise Linux 5 ships with GNOME 2.16, and Red Hat Enterprise Linux 6 ships with GNOME 2.28.

#### Table 2.1. Desktop Components Comparison

| Red Hat Enterprise Linux Version          |        |        |         |
|---|--------|--------|---------|
| GNOME Power Management Desktop Components | 4      | 5      | 6       |
| hal                                       | 0.4.2  | 0.5.8  | 0.5.14  |
| udisks                                    | N/A    | N/A    | 1.0.1   |
| glib2                                     | 2.4.7  | 2.12.3 | 2.22.5  |
| gtk2                                      | 2.4.13 | 2.10.4 | 2.18.9  |
| gnome-vfs2                                | 2.8.2  | 2.16.2 | 2.24.2  |
| libglade2                                 | 2.4.0  | 2.6.0  | 2.6.4   |
| libgnomecanvas                            | 2.8.0  | 2.14.0 | 2.26.0  |
| gnome-desktop                             | 2.8.0  | 2.16.0 | 2.28.2  |
| gnome-media                               | 2.8.0  | 2.16.1 | 2.29.91 |
| gnome-python2                             | 2.6.0  | 2.16.0 | 2.28.0  |
| libgnome                                  | 2.8.0  | 2.16.0 | 2.28.0  |
| libgnomeui                                | 2.8.0  | 2.16.0 | 2.24.1  |
| libgnomeprint22                           | 2.8.0  | 2.12.1 | N/A     |
| libgnomeprintui22                         | 2.8.0  | 2.12.1 | N/A     |
| gnome-session                             | 2.8.0  | 2.16.0 | 2.28.0  |
| gnome-power-manager                       | N/A    | 2.16.0 | 2.28.3  |
| gnome-applets                             | 2.8.0  | 2.16.0 | 2.28.0  |
| gnome-panel                               | 2.8.1  | 2.16.1 | 2.30.2  |

### 2.2.5.2. API Changes for glib

There are a number of API changes for glib between versions.

#### Version 2.4 to Version 2.12

Some of the differences in glib between version 2.4 and 2.12 (or between Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5) are:

- GOption (a command line option parser)
- GKeyFile (a key/ini file parser)
- GObject toggle references
- GMappedFile (a map wrapper)
- GSlice (a fast memory allocator)
- GBookmarkFile (a bookmark file parser)
- Base64 encoding support
- Native atomic ops on s390
- Updated Unicode support to 5
- Atomic reference counting for GObject

### **Version 2.12 to Version 2.22**

Some of the differences in glib between version 2.12 and 2.22 (or between Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6) are:

- GSequence (a list data structure that is implemented as a balanced tree)
- GRegex (a PCRE regex wrapper)
- Support for monotonic clocks
- XDG user dirs support
- GIO (a VFS library to replace gnome-vfs)
- GChecksum (support for hash algorithms such as MD5 and SHA-256)
- GTest (a test framework)
- Support for sockets and network IO in GIO
- GHashTable performance improvements
- GMarkup performance improvements

Documentation for glib, including indexes of new and deprecated APIs, is shipped in the glib2-devel package.

### **2.2.5.3. API Changes for GTK+**

There are a number of API changes for GTK+ between versions.

#### **Version 2.4 to Version 2.10**

Some of the differences in GTK+ between version 2.4 and 2.10 (or between Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5) are:

- GtkIconView
- GtkAboutDialog
- GtkCellView
- GtkFileChooserButton
- GtkMenuToolButton
- GtkAssistant
- GtkLinkButton
- GtkRecentChooser
- GtkCellRendererCombo
- GtkCellRendererProgress
- GtkCellRendererAccel
- GtkCellRendererSpin
- GtkStatusIcon
- Printing Support
- Notebook tab DND support
- Ellipsisation support in labels, progressbars and treeviews
- Support rotated text
- Improved themability

### **Version 2.10 to Version 2.18**

Some of the differences in GTK+ between version 2.10 and 2.18 (or between Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5) are:

- GtkScaleButton
- GtkVolumeButton
- GtkInfoBar
- GtkBuilder to replace libglade
- New tooltips API
- GtkMountOperation
- gtk\_show\_uri
- Scale marks
- Links in labels

- Support runtime font configuration changes
- Use GIO

Documentation for GTK+, including indexes of new and deprecated APIs, is shipped in the `gtk2-devel` package.

## 2.2.6. NSS Shared Databases

The NSS shared database format, introduced on NSS 3.12, is now available in Red Hat Enterprise Linux 6. This encompasses a number of new features and components to improve access and usability.

Included, is the NSS certificate and key database which are now `sqlite`-based and allow for concurrent access. The legacy `key3.db` and `cert8.db` are also replaced with new SQL databases called `key4.db` and `cert9.db`. These new databases will store PKCS #11 token objects, which are the same as what is currently stored in `cert8.db` and `key3.db`.

Having support for shared databases enables a system-wide NSS database. It resides in `/etc/pki/nssdb` where globally trusted CA certificates become accessible to all applications. The command `rv = NSS_InitReadWrite("sql:/etc/pki/nssdb");` initializes NSS for applications. If the application is run with root privileges, then the system-wide database is available on a read and write basis. However, if it is run with normal user privileges it becomes read only.

Additionally, a PEM PKCS #11 module for NSS allows applications to load into memory certificates and keys stored in PEM-formatted files (for example, those produced by `openssl`).

### 2.2.6.1. Backwards Compatibility

The binary compatibility guarantees made by NSS upstream are preserved in NSS for Red Hat Enterprise Linux 6. This guarantee states that NSS used in Red Hat Enterprise Linux 6 is backwards compatible with all older NSS 3.x shared libraries. Therefore, a program linked with an older NSS 3.x shared library will work without recompiling or relinking, and any applications that restrict the use of NSS APIs to the NSS Public Functions remain compatible with future versions of the NSS shared libraries.

### 2.2.6.2. NSS Shared Databases Documentation

Mozilla's wiki page explains the system-wide database rationale in great detail and can be accessed [http://wiki.mozilla.org/NSS\\_Shared\\_DB\\_And\\_LINUX](http://wiki.mozilla.org/NSS_Shared_DB_And_LINUX).

## 2.2.7. Python

The `python` package adds support for the Python programming language. This package provides the object and cached bytecode files required to enable runtime support for basic Python programs. It also contains the `python` interpreter and the `pydoc` documentation tool. The `python-devel` package contains the libraries and header files required for developing Python extensions.

Red Hat Enterprise Linux also ships with numerous `python`-related packages. By convention, the names of these packages have a `python` prefix or suffix. Such packages are either library extensions or python bindings to an existing library. For instance, `dbus-python` is a Python language binding for D-Bus.

Note that both cached bytecode (`*.pyc`/`*.pyo` files) and compiled extension modules (`*.so` files) are incompatible between Python 2.4 (used in Red Hat Enterprise Linux 5) and Python 2.6 (used in Red Hat Enterprise Linux 6). As such, you will be required to rebuild any extension modules you use that are not

part of Red Hat Enterprise Linux.

### 2.2.7.1. Python Updates

The Red Hat Enterprise Linux 6 version of Python features various language changes. For information about these changes, see the following project resources:

- What's New in Python 2.5: <http://docs.python.org/whatsnew/2.5.html>
- What's New in Python 2.6: <http://docs.python.org/whatsnew/2.6.html>

Both resources also contain advice on porting code developed using previous Python versions.

### 2.2.7.2. Python Documentation

For more information about Python, see `man python`. You can also install `python-docs`, which provides HTML manuals and references in the following location:

`file:///usr/share/doc/python-docs-version/html/index.html`

For details on library and language components, use `pydoc component_name`. For example, `pydoc math` will display the following information about the `math` Python module:

```

Help on module math:

NAME
  math

FILE
  /usr/lib64/python2.6/lib-dynload/mathmodule.so

DESCRIPTION
  This module is always available. It provides access to the
  mathematical functions defined by the C standard.

FUNCTIONS
  acos[...]
  acos(x)

  Return the arc cosine (measured in radians) of x.

  acosh[...]
  acosh(x)

  Return the hyperbolic arc cosine (measured in radians) of x.

  asin(...)
  asin(x)

  Return the arc sine (measured in radians) of x.

  asinh[...]
  asinh(x)

  Return the hyperbolic arc sine (measured in radians) of x.

```

The main site for the Python development project is hosted on [python.org](http://python.org).

## 2.2.8. Java

The `java-1.6.0-openjdk` package adds support for the Java programming language. This package provides the **java** interpreter. The `java-1.6.0-openjdk-devel` package contains the **javac** compiler, as well as the libraries and header files required for developing Java extensions.

Similarly, Red Hat Enterprise Linux also provides Java 7 via the `java-1.7.0-openjdk*` packages and Java 8 via the `java-1.8.0-openjdk*` packages.

### 2.2.8.1. Java Documentation

For more information about Java, see `man java`. Some associated utilities also have their own respective `man` pages.

You can also install other Java documentation packages for more details about specific Java utilities. By convention, such documentation packages have the **javadoc** suffix (for example, `dbus-java-javadoc`).

The main site for the development of Java is hosted on <http://openjdk.java.net/>. The main site for the library runtime of Java is hosted on <http://icedtea.classpath.org>.

## 2.2.9. Ruby

The **ruby** package provides the Ruby interpreter and adds support for the Ruby programming language. The **ruby-devel** package contains the libraries and header files required for developing Ruby extensions.

Red Hat Enterprise Linux also ships with numerous **ruby**-related packages. By convention, the names of these packages have a **ruby** or **rubym** prefix or suffix. Such packages are either library extensions or Ruby bindings to an existing library.

Examples of **ruby**-related packages include:

- `ruby-flexmock`
- `rubym-flexmock`
- `rubymgems`
- `ruby-irb`
- `ruby-libguestfs`
- `ruby-libs`
- `ruby-qpid`
- `ruby-rdoc`
- `ruby-ri`
- `ruby-saslwrapper`
- `ruby-static`



- `ruby-tcltk`

For information about updates to the Ruby language in Red Hat Enterprise Linux 6, see the following resources:

- `file:///usr/share/doc/ruby-version/NEWS`
- `file:///usr/share/doc/ruby-version/NEWS-version`

### 2.2.9.1. Ruby Documentation

For more information about Ruby, see `man ruby`. You can also install `ruby-docs`, which provides HTML manuals and references in the following location:

`file:///usr/share/doc/ruby-docs-version/`

The main site for the development of Ruby is hosted on <http://www.ruby-lang.org>. The <http://www.ruby-doc.org> site also contains Ruby documentation.

### 2.2.10. Perl

The `perl` package adds support for the Perl programming language. This package provides Perl core modules, the Perl Language Interpreter, and the PerlDoc tool.

Red Hat also provides various perl modules in package form; these packages are named with the `perl-*` prefix. These modules provide stand-alone applications, language extensions, Perl libraries, and external library bindings.

#### 2.2.10.1. Perl Updates

For a full list of the differences between the Perl versions see the following documents:

##### Perl 5.12 Updates

Perl 5.12 has the following updates:

- Perl conforms closer to the Unicode standard.
- Experimental APIs allow Perl to be extended with "pluggable" keywords and syntax.
- Perl will be able to keep accurate time well past the "Y2038" barrier.
- Package version numbers can be directly specified in "package" statements.
- Perl warns the user about the use of deprecated features by default.

The Perl 5.12 delta can be accessed at <http://perldoc.perl.org/perl5120delta.html>.

##### Perl 5.14 Updates

Perl 5.14 has the following updates:

- Unicode 6.0 support.
- Improved support for IPv6.
- Easier auto-configuration of the CPAN client.

- A new `/r` flag that makes `s///` substitutions non-destructive.
- New regular expression flags to control whether matched strings should be treated as ASCII or Unicode.
- New **package** `Foo { }` syntax.
- Less memory and CPU usage than previous releases.
- A number of bug fixes.

The Perl 5.14 delta can be accessed at <http://perldoc.perl.org/perl5140delta.html>.

### Perl 5.16 Updates

Perl 5.14 has the following updates:

- Support for Unicode 6.1.
- `$$` variable is writable.
- Improved debugger.
- Accessing Unicode database files directly is now deprecated; use `Unicode::UCD` instead.
- `Version::Requirements` is deprecated in favor of `CPAN::Meta::Requirements`.
- A number of perl4 libraries are removed:
  - `abbrev.pl`
  - `assert.pl`
  - `bigfloat.pl`
  - `bigint.pl`
  - `bigrat.pl`
  - `cacheout.pl`
  - `complete.pl`
  - `ctime.pl`
  - `dotsh.pl`
  - `exceptions.pl`
  - `fastcwd.pl`
  - `flush.pl`
  - `getcwd.pl`
  - `getopt.pl`
  - `getopts.pl`

- hostname.pl
- importenv.pl
- lib/find{,depth}.pl
- look.pl
- newgetopt.pl
- open2.pl
- open3.pl
- pwd.pl
- hellwords.pl
- stat.pl
- tainted.pl
- termcap.pl
- timelocal.pl

The Perl 5.16 delta can be accessed at <http://perldoc.perl.org/perl5160delta.html>.

### 2.2.10.2. Installation

Perl's capabilities can be extended by installing additional modules. These modules come in the following forms:

#### Official Red Hat RPM

The official module packages can be installed with **yum** or **rpm** from the Red Hat Enterprise Linux repositories. They are installed to **/usr/share/perl5** and either **/usr/lib/perl5** for 32bit architectures or **/usr/lib64/perl5** for 64bit architectures.

#### Modules from CPAN

Use the **cpan** tool provided by the perl-CPAN package to install modules directly from the CPAN website. They are installed to **/usr/local/share/perl5** and either **/usr/local/lib/perl5** for 32bit architectures or **/usr/local/lib64/perl5** for 64bit architectures.

#### Third party module package

Third party modules are installed to **/usr/share/perl5/vendor\_perl** and either **/usr/lib/perl5/vendor\_perl** for 32bit architectures or **/usr/lib64/perl5/vendor\_perl** for 64bit architectures.

#### Custom module package / manually installed module

These should be placed in the same directories as third-party modules. That is, **/usr/share/perl5/vendor\_perl** and either **/usr/lib/perl5/vendor\_perl** for 32bit architectures or **/usr/lib64/perl5/vendor\_perl** for 64bit architectures.

**WARNING**

If an official version of a module is already installed, installing its non-official version can create conflicts in the `/usr/share/man` directory.

**2.2.10.3. Perl Documentation**

The `perldoc` tool provides documentation on language and core modules. To learn more about a module, use `perldoc module_name`. For example, `perldoc CGI` will display the following information about the CGI core module:

```

NAME
  CGI - Handle Common Gateway Interface requests and responses

SYNOPSIS
  use CGI;

  my $q = CGI->new;

  [...]

DESCRIPTION
  CGI.pm is a stable, complete and mature solution for processing and
  preparing HTTP requests and responses. Major features including
  processing form submissions, file uploads, reading and writing cookies,
  query string generation and manipulation, and processing and preparing
  HTTP headers. Some HTML generation utilities are included as well.

  [...]

PROGRAMMING STYLE
  There are two styles of programming with CGI.pm, an object-oriented
  style and a function-oriented style. In the object-oriented style you
  create one or more CGI objects and then use object methods to create the
  various elements of the page. Each CGI object starts out with the list of
  named parameters that were passed to your CGI script by the server.

  [...]

```

For details on Perl functions, use `perldoc -f function_name`. For example, `perldoc -f split` will display the following information about the `split` function:

```

split /PATTERN/,EXPR,LIMIT
split /PATTERN/,EXPR
split /PATTERN/
split  Splits the string EXPR into a list of strings and returns
that list. By default, empty leading fields are preserved, and empty
trailing ones are deleted. (If all fields are empty, they are considered
to be trailing.)

```

In scalar context, returns the number of fields found. In scalar and void context it splits into the `@_` array. Use of `split` in scalar and void context is deprecated, however, because it clobbers your subroutine arguments.

If `EXPR` is omitted, splits the `$_` string. If `PATTERN` is also omitted, splits on whitespace (after skipping any leading whitespace). Anything matching `PATTERN` is taken to be a delimiter separating the fields. (Note that the delimiter may be longer than one character.)

[...]

Current PerlDoc documentation can be found on [perldoc.perl.org](http://perldoc.perl.org).

Core and external modules are documented on the [Comprehensive Perl Archive Network](#).

## CHAPTER 3. COMPILING AND BUILDING

Red Hat Enterprise Linux 6 includes many packages used for software development, including tools for compiling and building source code. This chapter discusses several of these packages and tools used to compile source code.

### 3.1. GNU COMPILER COLLECTION (GCC)

The GNU Compiler Collection (GCC) is a set of tools for compiling a variety of programming languages (including C, C++, ObjectiveC, ObjectiveC++, Fortran, and Ada) into highly optimized machine code. These tools include various compilers (like **gcc** and **g++**), run-time libraries (like **libgcc**, **libstdc++**, **libgfortran**, and **libgomp**), and miscellaneous other utilities.

#### 3.1.1. Language Compatibility

Application Binary Interfaces specified by the GNU C, C++, Fortran and Java Compiler include:

- Calling conventions. These specify how arguments are passed to functions and how results are returned from functions.
- Register usage conventions. These specify how processor registers are allocated and used.
- Object file formats. These specify the representation of binary object code.
- Size, layout, and alignment of data types. These specify how data is laid out in memory.
- Interfaces provided by the runtime environment. Where the documented semantics do not change from one version to another they must be kept available and use the same name at all times.

The default system C compiler included with Red Hat Enterprise Linux 6 is largely compatible with the C99 ABI standard. Deviations from the C99 standard in GCC 4.4 are tracked [online](#).

In addition to the C ABI, the Application Binary Interface for the GNU C++ Compiler specifies the binary interfaces required to support the C++ language, such as:

- Name mangling and demangling
- Creation and propagation of exceptions
- Formatting of run-time type information
- Constructors and destructors
- Layout, alignment, and padding of classes and derived classes
- Virtual function implementation details, such as the layout and alignment of virtual tables

The default system C++ compiler included with Red Hat Enterprise Linux 6 conforms to the C++ ABI defined by the [Itanium C++ ABI \(1.86\)](#).

Although every effort has been made to keep each version of GCC compatible with previous releases, some incompatibilities do exist.

#### **ABI incompatibilities between Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 5**

The following is a list of known incompatibilities between the Red Hat Enterprise Linux 6 and 5 toolchains.

- Passing/returning structs with flexible array members by value changed in some cases on Intel 64 and AMD64.
- Passing/returning of unions with long double members by value changed in some cases on Intel 64 and AMD64.
- Passing/returning structs with complex float member by value changed in some cases on Intel 64 and AMD64.
- Passing of 256-bit vectors on x86, Intel 64 and AMD64 platforms changed when `-mavx` is used.
- There have been multiple changes in passing of `_Decimal{32,64,128}` types and aggregates containing those by value on several targets.
- Packing of packed char bitfields changed in some cases.

### ABI incompatibilities between Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 4

The following is a list of known incompatibilities between the Red Hat Enterprise Linux 5 and 4 toolchains.

- There have been changes in the library interface specified by the C++ ABI for thread-safe initialization of function-scope static variables.
- On Intel 64 and AMD64, the medium model for building applications where data segment exceeds 4GB, was redesigned to match the latest ABI draft at the time. The ABI change results in incompatibility among medium model objects.

The compiler flag `-Wabi` can be used to get diagnostics indicating where these constructs appear in source code, though it will not catch every single case. This flag is especially useful for C++ code to warn whenever the compiler generates code that is known to be incompatible with the vendor-neutral C++ ABI.

Excluding the incompatibilities listed above, the GCC C and C++ language ABIs are *mostly* ABI compatible. The vast majority of source code will not encounter any of the known issues, and can be considered compatible.

Compatible ABIs allow the objects created by compiling source code to be portable to other systems. In particular, for Red Hat Enterprise Linux, this allows for *upward* compatibility. Upward compatibility is defined as the ability to link shared libraries and objects, created using a version of the compilers in a particular Red Hat Enterprise Linux release, with no problems. This includes new objects compiled on subsequent Red Hat Enterprise Linux releases.

The C ABI is considered to be stable, and has been so since at least Red Hat Enterprise Linux 3 (again, barring any incompatibilities mentioned in the above lists). Libraries built on Red Hat Enterprise Linux 3 and later can be linked to objects created on a subsequent environment (Red Hat Enterprise Linux 4, Red Hat Enterprise Linux 5, and Red Hat Enterprise Linux 6).

The C++ ABI is considered to be stable, but less stable than the C ABI, and only as of Red Hat Enterprise Linux 4 (corresponding to GCC version 3.4 and above.). As with C, this is only an upward compatibility. Libraries built on Red Hat Enterprise Linux 4 and above can be linked to objects created on a subsequent environment (Red Hat Enterprise Linux 5, and Red Hat Enterprise Linux 6).

To force GCC to generate code compatible with the C++ ABI in Red Hat Enterprise Linux releases prior to Red Hat Enterprise Linux 4, some developers have used the `-fabi-version=1` option. This practice

is not recommended. Objects created this way are indistinguishable from objects conforming to the current stable ABI, and can be linked (incorrectly) amongst the different ABIs, especially when using new compilers to generate code to be linked with old libraries that were built with tools prior to Red Hat Enterprise Linux 4.



### WARNING

The above incompatibilities make it incredibly difficult to maintain ABI shared library sanity between releases, especially when developing custom libraries with multiple dependencies outside of the core libraries. Therefore, if shared libraries are developed, it is *highly* recommend that a new version is built for each Red Hat Enterprise Linux release.

### 3.1.2. Object Compatibility and Interoperability

Two items that are important are the changes and enhancements in the underlying tools used by the compiler, and the compatibility between the different versions of a language's compiler.

Changes and new features in tools like **ld** (distributed as part of the **binutils** package) or in the dynamic loader (**ld.so**, distributed as part of the **glibc** package) can subtly change the object files that the compiler produces. These changes mean that object files moving to the current release of Red Hat Enterprise Linux from previous releases may lose functionality, behave differently at runtime, or otherwise interoperate in a diminished capacity. Known problem areas include:

- **ld --build-id**

In Red Hat Enterprise Linux 6 this is passed to **ld** by default, whereas Red Hat Enterprise Linux 5 **ld** doesn't recognize it.

- **as .cfi\_sections** support

In Red Hat Enterprise Linux 6 this directive allows **.debug\_frame**, **.eh\_frame** or both to be omitted from **.cfi\*** directives. In Red Hat Enterprise Linux 5 only **.eh\_frame** is omitted.

- **as, ld, ld.so, and gdb STB\_GNU\_UNIQUE** and **%gnu\_unique\_symbol** support

In Red Hat Enterprise Linux 6 more debug information is generated and stored in object files. This information relies on new features detailed in the **DWARF** standard, and also on new extensions not yet standardized. In Red Hat Enterprise Linux 5, tools like **as**, **ld**, **gdb**, **objdump**, and **readelf** may not be prepared for this new information and may fail to interoperate with objects created with the newer tools. In addition, Red Hat Enterprise Linux 5 produced object files do not support these new features; these object files may be handled by Red Hat Enterprise Linux 6 tools in a sub-optimal manner.

An outgrowth of this enhanced debug information is that the debuginfo packages that ship with system libraries allow you to do useful source level debugging into system libraries if they are installed. See [Section 4.2, “Installing Debuginfo Packages”](#) for more information on debuginfo packages.

Object file changes, such as the ones listed above, may interfere with the portable use of **prelink**.



### 3.1.3. Running GCC

To compile using GCC tools, first install the `binutils` and `gcc` packages. Doing so will also install several dependencies.

In brief, the tools work via the `gcc` command. This is the main driver for the compiler. It can be used from the command line to pre-process or compile a source file, link object files and libraries, or perform a combination thereof. By default, `gcc` takes care of the details and links in the provided `libgcc` library.

Conversely, using GCC tools from the command line interface consumes less system resources. This also allows finer-grained control over compilers; GCC's command line tools can even be used outside of the graphical mode (runlevel 5).

#### 3.1.3.1. Simple C Usage

Basic compilation of a C language program using GCC is easy. Start with the following simple program:

##### Example 3.1. `hello.c`

```
#include <stdio.h>
int main()
{
    printf ("Hello world!\n");
    return 0;
}
```

The following procedure illustrates the compilation process for C in its most basic form.

#### Procedure 3.1. Compiling a 'Hello World' C Program

1. Compile [Example 3.1](#), "`hello.c`" into an executable with:

```
~]$ gcc hello.c -o hello
```

Ensure that the resulting binary `hello` is in the same directory as `hello.c`.

2. Run the `hello` binary, that is, `./hello`.

#### 3.1.3.2. Simple C++ Usage

Basic compilation of a C++ language program using GCC is similar. Start with the following simple program:

##### Example 3.2. `hello.cc`

```
#include <iostream>
using namespace std;
int main()
{
    cout << "Hello World!" << endl;
    return 0;
}
```

The following procedure illustrates the compilation process for C++ in its most basic form.

### Procedure 3.2. Compiling a 'Hello World' C++ Program

1. Compile [Example 3.2](#), “`hello.cc`” into an executable with:

```
~]$ g++ hello.cc -o hello
```

Ensure that the resulting binary `hello` is in the same directory as `hello.cc`.

2. Run the `hello` binary, that is, `./hello`.

### 3.1.3.3. Simple Multi-File Usage

To use basic compilation involving multiple files or object files, start with the following two source files:

#### Example 3.3. `one.c`

```
#include <stdio.h>
void hello()
{
    printf("Hello world!\n");
}
```

#### Example 3.4. `two.c`

```
extern void hello();
int main()
{
    hello();
    return 0;
}
```

The following procedure illustrates a simple, multi-file compilation process in its most basic form.

### Procedure 3.3. Compiling a Program with Multiple Source Files

1. Compile [Example 3.3](#), “`one.c`” into an executable with:

```
~]$ gcc -c one.c -o one.o
```

Ensure that the resulting binary `one.o` is in the same directory as `one.c`.

2. Compile [Example 3.4](#), “`two.c`” into an executable with:

```
~]$ gcc -c two.c -o two.o
```

Ensure that the resulting binary `two.o` is in the same directory as `two.c`.

3. Compile the two object files `one.o` and `two.o` into a single executable with:

```
~]$ gcc one.o two.o -o hello
```

Ensure that the resulting binary `hello` is in the same directory as `one.o` and `two.o`.

4. Run the `hello` binary, that is, `./hello`.

### 3.1.3.4. Recommended Optimization Options

Different projects require different optimization options. There is no one-size-fits-all approach when it comes to optimization, but here are a few guidelines to keep in mind.

#### Instruction selection and tuning

It is very important to choose the correct architecture for instruction scheduling. By default GCC produces code optimized for the most common processors, but if the CPU on which your code will run is known, the corresponding `-mtune=` option to optimize the instruction scheduling, and `-march=` option to optimize the instruction selection should be used.

The option `-mtune=` optimizes instruction scheduling to fit your architecture by tuning everything except the ABI and the available instruction set. This option will not choose particular instructions, but instead will tune your program in such a way that executing on a particular architecture will be optimized. For example, if an Intel Core2 CPU will predominantly be used, choose `-mtune=core2`. If the wrong choice is made, the program will still run, but not optimally on the given architecture. The architecture on which the program will most likely run should always be chosen.

The option `-march=` optimizes instruction selection. As such, it is important to choose correctly as choosing incorrectly will cause your program to fail. This option selects the instruction set used when generating code. For example, if the program will be run on an AMD K8 core based CPU, choose `-march=k8`. Specifying the architecture with this option will imply `-mtune=`.

The `-mtune=` and `-march=` commands should only be used for tuning and selecting instructions within a given architecture, not to generate code for a different architecture (also known as cross-compiling). For example, this is not to be used to generate PowerPC code from an Intel 64 and AMD64 platform.

For a complete list of the available options for both `-march=` and `-mtune=`, see the GCC documentation available here: [GCC 4.4.4 Manual: Hardware Models and Configurations](#)

#### General purpose optimization flags

The compiler flag `-O2` is a good middle of the road option to generate fast code. It produces the best optimized code when the resulting code size is not large. Use this when unsure what would best suit.

When code size is not an issue, `-O3` is preferable. This option produces code that is slightly larger but runs faster because of a more frequent inline of functions. This is ideal for floating point intensive code.

The other general purpose optimization flag is `-Os`. This flag also optimizes for size, and produces faster code in situations where a smaller footprint will increase code locality, thereby reducing cache misses.

Use `-frecord-gcc-switches` when compiling objects. This records the options used to build objects into objects themselves. After an object is built, it determines which set of options were used to build it. The set of options are then recorded in a section called `.GCC.command.line` within the object and can be examined with the following:

```
$ gcc -frecord-gcc-switches -O3 -Wall hello.c -o hello
```

```
$ readelf --string-dump=.GCC.command.line hello
String dump of section '.GCC.command.line':
 [   0]  hello.c
 [   8]  -mtune=generic
 [  17]  -O3
 [  1b]  -Wall
 [  21]  -frecord-gcc-switches
```

It is very important to test and try different options with a representative data set. Often, different modules or objects can be compiled with different optimization flags in order to produce optimal results. See [Section 3.1.3.5, “Using Profile Feedback to Tune Optimization Heuristics”](#) for additional optimization tuning.

### 3.1.3.5. Using Profile Feedback to Tune Optimization Heuristics

During the transformation of a typical set of source code into an executable, tens of hundreds of choices must be made about the importance of speed in one part of code over another, or code size as opposed to code speed. By default, these choices are made by the compiler using reasonable heuristics, tuned over time to produce the optimum runtime performance. However, GCC also has a way to teach the compiler to optimize executables for a specific machine in a specific production environment. This feature is called profile feedback.

Profile feedback is used to tune optimizations such as:

- Inlining
- Branch prediction
- Instruction scheduling
- Inter-procedural constant propagation
- Determining of hot or cold functions

Profile feedback compiles a program first to generate a program that is run and analyzed and then a second time to optimize with the gathered data.

#### Procedure 3.4. Using Profile Feedback

1. The application must be instrumented to produce profiling information by compiling it with **-fprofile-generate**.
2. Run the application to accumulate and save the profiling information.
3. Recompile the application with **-fprofile-use**.

Step three will use the profile information gathered in step one to tune the compiler's heuristics while optimizing the code into a final executable.

#### Procedure 3.5. Compiling a Program with Profiling Feedback

1. Compile **source.c** to include profiling instrumentation:  
**gcc source.c -fprofile-generate -O2 -o executable**
2. Run **executable** to gather profiling information:

`./executable`

3. Recompile and optimize `source.c` with profiling information gathered in step one:

```
gcc source.c -fprofile-use -O2 -o executable
```

Multiple data collection runs, as seen in step two, will accumulate data into the profiling file instead of replacing it. This allows the executable in step two to be run multiple times with additional representative data in order to collect even more information.

The executable must run with representative levels of both the machine being used and a respective data set large enough for the input required. This ensures optimal results are achieved.

By default, GCC will generate the profile data into the directory where step one was performed. To generate this information elsewhere, compile with `-fprofile-dir=DIR` where **DIR** is the preferred output directory.



### WARNING

The format of the compiler feedback data file changes between compiler versions. It is imperative that the program compilation is repeated with each version of the compiler.

#### 3.1.3.6. Using 32-bit compilers on a 64-bit host

On a 64-bit host, GCC will build executables that can only run on 64-bit hosts. However, GCC can be used to build executables that will run both on 64-bit hosts and on 32-bit hosts.

To build 32-bit binaries on a 64-bit host, first install 32-bit versions of any supporting libraries the executable may require. This must at least include supporting libraries for **glibc** and **libgcc**, and **libstdc++** if the program is a C++ program. On Intel 64 and AMD64, this can be done with:

```
yum install glibc-devel.i686 libgcc.i686 libstdc++-devel.i686
```

There may be cases where it is useful to to install additional 32-bit libraries that a program may require. For example, if a program uses the **db4-devel** libraries to build, the 32-bit version of these libraries can be installed with:

```
yum install db4-devel.i686
```



### NOTE

The **.i686** suffix on the x86 platform (as opposed to **x86-64**) specifies a 32-bit version of the given package. For PowerPC architectures, the suffix is **ppc** (as opposed to **ppc64**).

After the 32-bit libraries have been installed, the `-m32` option can be passed to the compiler and linker to produce 32-bit executables. Provided the supporting 32-bit libraries are installed on the 64-bit system, this executable will be able to run on both 32-bit systems and 64-bit systems.

**Procedure 3.6. Compiling a 32-bit Program on a 64-bit Host**

1. On a 64-bit system, compile `hello.c` into a 64-bit executable with:

```
gcc hello.c -o hello64
```

2. Ensure that the resulting executable is a 64-bit binary:

```
$ file hello64
hello64: ELF 64-bit LSB executable, x86-64, version 1
(GNU/Linux), dynamically linked (uses shared libs), for GNU/Linux
2.6.18, not stripped
$ ldd hello64
linux-vdso.so.1 => (0x00007fff242dd000)
libc.so.6 => /lib64/libc.so.6 (0x00007f0721514000)
/lib64/ld-linux-x86-64.so.2 (0x00007f0721893000)
```

The command `file` on a 64-bit executable will include **ELF 64-bit** in its output, and `ldd` will list `/lib64/libc.so.6` as the main C library linked.

3. On a 64-bit system, compile `hello.c` into a 32-bit executable with:

```
gcc -m32 hello.c -o hello32
```

4. Ensure that the resulting executable is a 32-bit binary:

```
$ file hello32
hello32: ELF 32-bit LSB executable, Intel 80386,
version 1 (GNU/Linux), dynamically linked (uses shared libs), for
GNU/Linux 2.6.18, not stripped
$ ldd hello32
linux-gate.so.1 => (0x007eb000)
libc.so.6 => /lib/libc.so.6 (0x00b13000)
/lib/ld-linux.so.2 (0x00cd7000)
```

The command `file` on a 32-bit executable will include **ELF 32-bit** in its output, and `ldd` will list `/lib/libc.so.6` as the main C library linked.

If you have not installed the 32-bit supporting libraries you will get an error similar to this for C code:

```
$ gcc -m32 hello32.c -o hello32
/usr/bin/ld: crt1.o: No such file: No such file or directory
collect2: ld returned 1 exit status
```

A similar error would be triggered on C++ code:

```
$ g++ -m32 hello32.cc -o hello32-c++
In file included from /usr/include/features.h:385,
    from /usr/lib/gcc/x86_64-redhat-
linux/4.4.4/../../../../include/c++/4.4.4/x86_64-redhat-
linux/32/bits/os_defines.h:39,
    from /usr/lib/gcc/x86_64-redhat-
linux/4.4.4/../../../../include/c++/4.4.4/x86_64-redhat-
linux/32/bits/c++config.h:243,
```

```
    from /usr/lib/gcc/x86_64-redhat-  
linux/4.4.4/../../../../include/c++/4.4.4/iostream:39,  
    from hello32.cc:1:  
/usr/include/gnu/stubs.h:7:27: error: gnu/stubs-32.h: No such file or  
directory
```

These errors indicate that the supporting 32-bit libraries have not been properly installed as explained at the beginning of this section.

Also important is to note that building with `-m32` will not adapt or convert a program to resolve any issues arising from 32/64-bit incompatibilities. For tips on writing portable code and converting from 32-bits to 64-bits, see the paper entitled *Porting to 64-bit GNU/Linux Systems* in the [Proceedings of the 2003 GCC Developers Summit](#).

### 3.1.4. GCC Documentation

For more information about GCC compilers, see the `man` pages for `cpp`, `gcc`, `g++`, `gcj`, and `gfortran`.

The following online user manuals are also available:

- [GCC 4.4.4 Manual](#)
- [GCC 4.4.4 GNU Fortran Manual](#)
- [GCC 4.4.4 GCJ Manual](#)
- [GCC 4.4.4 CPP Manual](#)
- [GCC 4.4.4 GNAT Reference Manual](#)
- [GCC 4.4.4 GNAT User's Guide](#)
- [GCC 4.4.4 GNU OpenMP Manual](#)

The main site for the development of GCC is [gcc.gnu.org](http://gcc.gnu.org).

## 3.2. AUTOTOOLS

GNU Autotools is a suite of command line tools that allow developers to build applications on different systems, regardless of the installed packages or even Linux distribution. These tools aid developers in creating a **configure** script. This script runs prior to builds and creates the top-level **Makefiles** required to build the application. The **configure** script may perform tests on the current system, create additional files, or run other directives as per parameters provided by the builder.

The Autotools suite's most commonly-used tools are:

### **autoconf**

Generates the **configure** script from an input file (**configure.ac**, for example)

### **automake**

Creates the **Makefile** for a project on a specific system

### **autoscan**

Generates a preliminary input file (that is, **configure.scan**), which can be edited to create a final **configure.ac** to be used by **autoconf**

All tools in the Autotools suite are part of the **Development Tools** group package. You can install this package group to install the entire Autotools suite, or use **yum** to install any tools in the suite as you wish.

### 3.2.1. Autotools Plug-in for Eclipse

The Autotools suite is also integrated into the Eclipse IDE via the Autotools plug-in. This plug-in provides an Eclipse graphical user interface for Autotools, which is suitable for most C/C++ projects.

As of Red Hat Enterprise Linux 6, this plug-in only supports two templates for new C/C++ projects:

- An empty project
- A "hello world" application

The empty project template is used when importing projects into the C/C++ Development Toolkit that already support Autotools. Future updates to the Autotools plug-in will include additional graphical user interfaces (wizards, for example) for creating shared libraries and other complex scenarios.

The Red Hat Enterprise Linux 6 version of the Autotools plug-in also does not integrate **git** or **mercurial** into Eclipse. As such, Autotools projects that use **git** repositories will be required to be checked out *outside* the Eclipse workspace. Afterwards, you can specify the source location for such projects in Eclipse. Any repository manipulation (commits, or updates for example) are done via the command line.

### 3.2.2. Configuration Script

The most crucial function of Autotools is the creation of the **configure** script. This script tests systems for tools, input files, and other features it can use in order to build the project <sup>[1]</sup>. The **configure** script generates a **Makefile** which allows the **make** tool to build the project based on the system configuration.

To create the **configure** script, first create an input file. Then feed it to an Autotools utility in order to create the **configure** script. This input file is typically **configure.ac** or **Makefile.am**; the former is usually processed by **autoconf**, while the later is fed to **automake**.

If a **Makefile.am** input file is available, the **automake** utility creates a **Makefile** template (that is, **Makefile.in**), which may see information collected at configuration time. For example, the **Makefile** may have to link to a particular library *if and only if* that library is already installed. When the **configure** script runs, **automake** will use the **Makefile.in** templates to create a **Makefile**.

If a **configure.ac** file is available instead, then **autoconf** will automatically create the **configure** script based on the macros invoked by **configure.ac**. To create a preliminary **configure.ac**, use the **autoscan** utility and edit the file accordingly.

### 3.2.3. Autotools Documentation

Red Hat Enterprise Linux 6 includes **man** pages for **autoconf**, **automake**, **autoscan** and most tools included in the Autotools suite. In addition, the Autotools community provides extensive documentation on **autoconf** and **automake** on the following websites:



- <http://www.gnu.org/software/autoconf/manual/autoconf.html>
- <http://www.gnu.org/software/autoconf/manual/automake.html>

The following is an online book describing the use of Autotools. Although the above online documentation is the recommended and most up to date information on Autotools, this book is a good alternative and introduction.

- <http://sourceware.org/autobook/>

For information on how to create Autotools input files, see:

- <http://www.gnu.org/software/autoconf/manual/autoconf.html#Making-configure-Scripts>
- <http://www.gnu.org/software/autoconf/manual/automake.html#Invoking-Automake>

The following upstream example also illustrates the use of Autotools in a simple **hello** program:

- <http://www.gnu.org/software/hello/manual/hello.html>

### 3.3. BUILD-ID UNIQUE IDENTIFICATION OF BINARIES

Each executable or shared library built with Red Hat Enterprise Linux Server 6 or later is assigned a unique identification 160-bit SHA-1 string, generated as a checksum of selected parts of the binary. This allows two builds of the same program on the same host to always produce consistent build-ids and binary content.

Display the build-id of a binary with the following command:

```
$ eu-readelf -n /bin/bash
[...]
Note section [ 3] '.note.gnu.build-id' of 36 bytes at offset 0x274:
  Owner  Data size  Type
  GNU    20      GNU_BUILD_ID
  Build ID: efdd0b5e69b0742fa5e5bad0771df4d1df2459d1
```

Unique identifiers of binaries are useful in cases such as analysing core files, documented [Section 4.2.1, “Installing Debuginfo Packages for Core Files Analysis”](#).

[1] For information about tests that **configure** can perform, see the following link:

<http://www.gnu.org/software/autoconf/manual/autoconf.html#Existing-Tests>

## CHAPTER 4. DEBUGGING


Useful, well-written software generally goes through several different phases of application development, allowing ample opportunity for mistakes to be made. Some phases come with their own set of mechanisms to detect errors. For example, during compilation an elementary semantic analysis is often performed to make sure objects, such as variables and functions, are adequately described.

The error-checking mechanisms performed during each application development phase aims to catch simple and obvious mistakes in code. The debugging phase helps to bring more subtle errors to light that fell through the cracks during routine code inspection.

### 4.1. ELF EXECUTABLE BINARIES

Red Hat Enterprise Linux uses ELF for executable binaries, shared libraries, or debuginfo files. Within these debuginfo ELF files, the DWARF format is used. Version 3 of DWARF is used in ELF files (that is, `gcc -g` is equivalent to `gcc -gdwarf-3`). DWARF debuginfo includes:

- names of all the compiled functions and variables, including their target addresses in binaries
- source files used for compilation, including their source line numbers
- local variables location



#### IMPORTANT

STABS is occasionally used with UNIX. STABS is an older, less capable format. Its use is discouraged by Red Hat. GCC and GDB support STABS production and consumption on a best effort basis only.

Within these ELF files, the GCC debuginfo level is also used. The default is level 2, where macro information is not present; level 3 has C/C++ macro definitions included, but the debuginfo can be very large with this setting. The command for the default `gcc -g` is the same as `gcc -g2`. To change the macro information to level three, use `gcc -g3`.

There are multiple levels of debuginfo available. Use the command `readelf -WS file` to see which sections are used in a file.

**Table 4.1. debuginfo levels**

| Binary State | Command  | Notes   |
|--------------|--|---|
| Stripped     | <code>strip file</code><br>or<br><code>gcc -s -o file</code> | Only the symbols required for runtime linkage with shared libraries are present.<br><br>ELF section in use: <code>.dynsym</code>                    |
| ELF symbols  | <code>gcc -o file</code>                                     | Only the names of functions and variables are present, no binding to the source files and no types.<br><br>ELF section in use: <code>.symtab</code> |

| Binary State                | Command                | Notes   |
|-----------------------------|------------------------|---|
| DWARF debuginfo with macros | <b>gcc -g -o file</b>  | The source file names and line numbers are known, including types.<br><br>ELF section in use: <b>.debug_*</b> |
| DWARF debuginfo with macros | <b>gcc -g3 -o file</b> | Similar to <b>gcc -g</b> but the macros are known to GDB.<br><br>ELF section in use:<br><b>.debug_macro</b>   |

**NOTE**

GDB never interprets the source files, it only displays them as text. Use **gcc -g** and its variants to store the information into DWARF.

Compiling a program or library with **gcc -rdynamic** is discouraged. For specific symbols, use **gcc -Wl, --dynamic-list=...** instead. If **gcc -rdynamic** is used, the **strip** command or **-s gcc** option have no effect. This is because all ELF symbols are kept in the binary for possible runtime linkage with shared libraries.

ELF symbols can be read by the **readelf -s file** command.

DWARF symbols are read by the **readelf -w file** command.

The command **readelf -wi file** is a good verification of debuginfo, compiled within your program. The commands **strip file** or **gcc -s** are commonly accidentally executed on the output during various compilation stages of the program.

The **readelf -w file** command can also be used to show a special section called **.eh\_frame** with a format and purpose is similar to the DWARF section **.debug\_frame**. The **.eh\_frame** section is used for runtime C++ exception resolution and is present even if **-g gcc** option was not used. It is kept in the primary RPM and is never present in the debuginfo RPMs.

Debuginfo RPMs contain the sections **.symtab** and **.debug\_\***. Neither **.eh\_frame**, **.eh\_frame\_hdr**, nor **.dynsym** are moved or present in debuginfo RPMs as those sections are needed during program runtime.

## 4.2. INSTALLING DEBUGINFO PACKAGES

Red Hat Enterprise Linux also provides **-debuginfo** packages for all architecture-dependent RPMs included in the operating system. A **packagename-debuginfo-version-release.architecture.rpm** package contains detailed information about the relationship of the package source files and the final installed binary. The debuginfo packages contain both **.debug** files, which in turn contain DWARF debuginfo and the source files used for compiling the binary packages.



## NOTE

Most of the debugger functionality is missed if attempting to debug a package without having its debuginfo equivalent installed. For example, the names of exported shared library functions will still be available, but the matching source file lines will not be without the debuginfo package installed.

Use **gcc** compilation option **-g** for your own programs. The debugging experience is better if no optimizations (gcc option **-O**, such as **-O2**) is applied with **-g**.

For Red Hat Enterprise Linux 6, the debuginfo packages are now available on a new channel on the Red Hat Network. To install the **-debuginfo** package of a package (that is, typically **packagename-debuginfo**), first the machine has to be subscribed to the corresponding Debuginfo channel. For example, for Red Hat Enterprise Server 6, the corresponding channel would be **Red Hat Enterprise Linux Server Debuginfo (v. 6)**.

Red Hat Enterprise Linux system packages are compiled with optimizations (gcc option **-O2**). This means that some variables will be displayed as **<optimized out>**. Stepping through code will 'jump' a little but a crash can still be analyzed. If some debugging information is missing because of the optimizations, the right variable information can be found by disassembling the code and matching it to the source manually. This is applicable only in exceptional cases and is not suitable for regular debugging.

For system packages, GDB informs the user if it is missing some debuginfo packages that limit its functionality.

```
gdb ls
[...]
Reading symbols from /bin/ls...(no debugging symbols found)...done.
Missing separate debuginfos, use: debuginfo-install coreutils-8.4-
16.el6.x86_64
(gdb) q
```

If the system package to be debugged is known, use the command suggested by GDB above. It will also automatically install all the debug packages *packagename* depends on.

```
# debuginfo-install packagename
```

### 4.2.1. Installing Debuginfo Packages for Core Files Analysis

A core file is a representation of the memory image at the time of a process crash. For bug reporting of system program crashes, Red Hat recommends the use of the ABRT tool, explained in the *Automatic Bug Reporting Tool* chapter in the *Red Hat Deployment Guide*. If ABRT is not suitable for your purposes, the steps it automates are explained here.

If the **ulimit -c unlimited** setting is in use when a process crashes, the core file is dumped into the current directory. The core file contains only the memory areas modified by the process from the original state of disk files. In order to perform a full analysis of a crash, a core file is required to have:

- the core file itself
- the executable binary which has crashed, such as **/usr/sbin/sendmail**
- all the shared libraries loaded in the binary when it crashed

- .debug files and source files (both stored in debuginfo RPMs) for the executable and all of its loaded libraries

For a proper analysis, either the exact **version-release.architecture** for all the RPMs involved or the same build of your own compiled binaries is needed. At the time of the crash, the application may have already recompiled or been updated by **yum** on the disk, rendering the files inappropriate for the core file analysis.

The core file contains build-ids of all the binaries involved. For more information on build-id, see [Section 3.3, “build-id Unique Identification of Binaries”](#). The contents of the core file can be displayed by:

```
$ eu-unstrip -n --core=./core.9814
0x400000+0x207000 2818b2009547f780a5639c904cded443e564973e@0x400284
/bin/sleep /usr/lib/debug/bin/sleep.debug [exe]
0x7fff26fff000+0x1000
1e2a683b7d877576970e4275d41a6aaec280795e@0x7fff26fff340 . - linux-
vdso.so.1
0x35e7e00000+0x3b6000
374add1ead31ccb449779bc7ee7877de3377e5ad@0x35e7e00280 /lib64/libc-
2.14.90.so /usr/lib/debug/lib64/libc-2.14.90.so.debug libc.so.6
0x35e7a00000+0x224000
3ed9e61c2b7e707ce244816335776afa2ad0307d@0x35e7a001d8 /lib64/ld-2.14.90.so
/usr/lib/debug/lib64/ld-2.14.90.so.debug ld-linux-x86-64.so.2
```

The meaning of the columns in each line are:

- The in-memory address where the specific binary was mapped to (for example, **0x400000** in the first line).
- The size of the binary (for example, **+0x207000** in the first line).
- The 160-bit SHA-1 build-id of the binary (for example, **2818b2009547f780a5639c904cded443e564973e** in the first line).
- The in-memory address where the build-id bytes were stored (for example, **@0x400284** in the first line).
- The on-disk binary file, if available (for example, **/bin/sleep** in the first line). This was found by **eu-unstrip** for this module.
- The on-disk debuginfo file, if available (for example, **/usr/lib/debug/bin/sleep.debug**). However, best practice is to use the binary file reference instead.
- The shared library name as stored in the shared library list in the core file (for example, **libc.so.6** in the third line).

For each build-id (for example, **ab/cdef0123456789012345678901234567890123**) a symbolic link is included in its debuginfo RPM. Using the **/bin/sleep** executable above as an example, the **coreutils-debuginfo** RPM contains, among other files:

```
lrwxrwxrwx 1 root root 24 Nov 29 17:07 /usr/lib/debug/.build-
id/28/18b2009547f780a5639c904cded443e564973e -> ../../../../../../bin/sleep*
lrwxrwxrwx 1 root root 21 Nov 29 17:07 /usr/lib/debug/.build-
id/28/18b2009547f780a5639c904cded443e564973e.debug ->
../../../../bin/sleep.debug
```

In some cases (such as loading a core file), GDB does not know the name, version, or release of a **`name-debuginfo-version-release.rpm`** package; it only knows the build-id. In such cases, GDB suggests a different command:

```
gdb -c ./core
[...]
Missing separate debuginfo for the main executable filename
Try: yum --disablerepo='*' --enablerepo='*debug*' install
/usr/lib/debug/.build-id/ef/dd0b5e69b0742fa5e5bad0771df4d1df2459d1
```

The *version-release.architecture* of the binary package *packagename-debuginfo-version-release.architecture.rpm* must be an exact match. If it differs then GDB cannot use the debuginfo package. Even the same *version-release.architecture* from a different build leads to an incompatible debuginfo package. If GDB reports a missing debuginfo, ensure to recheck:

### **rpm -q packagename packagename-debuginfo**

The *version-release.architecture* definitions should match.

### **rpm -V packagename packagename-debuginfo**

This command should produce no output, except possibly modified configuration files of *packagename*, for example.

### **rpm -qi packagename packagename-debuginfo**

The *version-release.architecture* should display matching information for Vendor, Build Date, and Build Host. For example, using a CentOS debuginfo RPM for a Red Hat Enterprise Linux RPM package will not work.

If the required build-id is known, the following command can query which RPM contains it:

```
$ repoquery --disablerepo='*' --enablerepo='*-debug*' -qf
/usr/lib/debug/.build-id/ef/dd0b5e69b0742fa5e5bad0771df4d1df2459d1
```

For example, a version of an executable which matches the core file can be installed by:

```
# yum --enablerepo='*-debug*' install $(eu-unstrip -n --core=./core.9814 |
sed -e 's#[^ ]* \(\.\.\)\([^\@ ]*\).*$/usr/lib/debug/.build-id/\1/\2#p' -e
's$/ .debug/')
```

Similar methods are available if the binaries are not packaged into RPMs and stored in yum repositories. It is possible to create local repositories with custom application builds by using **`/usr/bin/createrepo`**.

## 4.3. GDB

Fundamentally, like most debuggers, GDB manages the execution of compiled code in a very closely controlled environment. This environment makes possible the following fundamental mechanisms necessary to the operation of GDB:

- Inspect and modify memory within the code being debugged (for example, reading and setting variables).

- Control the execution state of the code being debugged, principally whether it's running or stopped.
- Detect the execution of particular sections of code (for example, stop running code when it reaches a specified area of interest to the programmer).
- Detect access to particular areas of memory (for example, stop running code when it accesses a specified variable).
- Execute portions of code (from an otherwise stopped program) in a controlled manner.
- Detect various programmatic asynchronous events such as signals.

The operation of these mechanisms rely mostly on information produced by a compiler. For example, to view the value of a variable, GDB has to know:

- The location of the variable in memory
- The nature of the variable

This means that displaying a double-precision floating point value requires a very different process from displaying a string of characters. For something complex like a structure, GDB has to know not only the characteristics of each individual elements in the structure, but the morphology of the structure as well.

GDB requires the following items in order to fully function:

### Debug Information

Much of GDB's operations rely on a program's *debug information*. While this information generally comes from compilers, much of it is necessary only while debugging a program, that is, it is not used during the program's normal execution. For this reason, compilers do not always make that information available by default — GCC, for instance, must be explicitly instructed to provide this debugging information with the `-g` flag.

To make full use of GDB's capabilities, it is *highly advisable* to make the debug information available first to GDB. GDB can only be of *very limited* use when run against code with no available debug information.

### Source Code

One of the most useful features of GDB (or any other debugger) is the ability to associate events and circumstances in program execution with their corresponding location in source code. This location normally refers to a specific line or series of lines in a source file. This, of course, would require that a program's source code be available to GDB at debug time.

#### 4.3.1. Simple GDB

GDB literally contains dozens of commands. This section describes the most fundamental ones.

##### **br (breakpoint)**

The breakpoint command instructs GDB to halt execution upon reaching a specified point in the execution. That point can be specified a number of ways, but the most common are just as the line number in the source file, or the name of a function. Any number of breakpoints can be in effect simultaneously. This is frequently the first command issued after starting GDB.

##### **r (run)**

The **run** command starts the execution of the program. If **run** is executed with any arguments, those arguments are passed on to the executable as if the program has been started normally. Users normally issue this command after setting breakpoints.

Before an executable is started, or once the executable stops at, for example, a breakpoint, the state of many aspects of the program can be inspected. The following commands are a few of the more common ways things can be examined.

### **p (print)**

The **print** command displays the value of the argument given, and that argument can be almost anything relevant to the program. Usually, the argument is the name of a variable of any complexity, from a simple single value to a structure. An argument can also be an expression valid in the current language, including the use of program variables and library functions, or functions defined in the program being tested.

### **bt (backtrace)**

The **backtrace** displays the chain of function calls used up until the execution was terminated. This is useful for investigating serious bugs (such as segmentation faults) with elusive causes.

### **l (list)**

When execution is stopped, the **list** command shows the line in the source code corresponding to where the program stopped.

The execution of a stopped program can be resumed in a number of ways. The following are the most common.

### **c (continue)**

The **continue** command restarts the execution of the program, which will continue to execute until it encounters a breakpoint, runs into a specified or emergent condition (for example, an error), or terminates.

### **n (next)**

Like **continue**, the **next** command also restarts execution; however, in addition to the stopping conditions implicit in the **continue** command, **next** will also halt execution at the next sequential line of code in the current source file.

### **s (step)**

Like **next**, the **step** command also halts execution at each sequential line of code in the current source file. However, if execution is currently stopped at a source line containing a *function call*, GDB stops execution after entering the function call (rather than executing it).

### **fini (finish)**

Like the aforementioned commands, the **finish** command resumes executions, but halts when execution returns from a function.

Finally, two essential commands:

### **q (quit)**



This terminates the execution.

## h (help)

The **help** command provides access to its extensive internal documentation. The command takes arguments: **help breakpoint** (or **h br**), for example, shows a detailed description of the **breakpoint** command. See the **help** output of each command for more detailed information.

## 4.3.2. Running GDB

This section will describe a basic execution of GDB, using the following simple program:

### hello.c

```
#include <stdio.h>

char hello[] = { "Hello, World!" };

int
main()
{
    fprintf (stdout, "%s\n", hello);
    return (0);
}
```

The following procedure illustrates the debugging process in its most basic form.

### Procedure 4.1. Debugging a 'Hello World' Program

1. Compile `hello.c` into an executable with the debug flag set, as in:

```
gcc -g -o hello hello.c
```

Ensure that the resulting binary **hello** is in the same directory as **hello.c**.

2. Run **gdb** on the **hello** binary, that is, **gdb hello**.
3. After several introductory comments, **gdb** will display the default GDB prompt:

```
(gdb)
```

4. The variable **hello** is global, so it can be seen even before the **main** procedure starts:

```
(gdb) p hello
$1 = "Hello, World!"
(gdb) p hello[0]
$2 = 72 'H'
(gdb) p *hello
$3 = 72 'H'
(gdb)
```

Note that the **print** targets **hello[0]** and **\*hello** require the evaluation of an expression, as does, for example, **\*(hello + 1)**:

```
(gdb) p *(hello + 1)
$4 = 101 'e'
```

5. Next, list the source:

```
(gdb) l
1      #include <stdio.h>
2
3      char hello[] = { "Hello, World!" };
4
5      int
6      main()
7      {
8          fprintf (stdout, "%s\n", hello);
9          return (0);
10     }
```

The **list** reveals that the **fprintf** call is on line 8. Apply a breakpoint on that line and resume the code:

```
(gdb) br 8
Breakpoint 1 at 0x80483ed: file hello.c, line 8.
(gdb) r
Starting program: /home/moller/tinkering/gdb-manual/hello

Breakpoint 1, main () at hello.c:8
8          fprintf (stdout, "%s\n", hello);
```

6. Finally, use the **next** command to step past the **fprintf** call, executing it:

```
(gdb) n
Hello, World!
9          return (0);
```

The following sections describe more complex applications of GDB.

### 4.3.3. Conditional Breakpoints

In many real-world cases, a program may perform its task well during the first few thousand times; it may then start crashing or encountering errors during its eight thousandth iteration of the task. Debugging programs like this can be difficult, as it is hard to imagine a programmer with the patience to issue a **continue** command thousands of times just to get to the iteration that crashed.

Situations like this are common in real life, which is why GDB allows programmers to attach conditions to a breakpoint. For example, consider the following program:

#### simple.c

```
#include <stdio.h>

main()
{
    int i;
```

```

    for (i = 0;; i++) {
    fprintf (stdout, "i = %d\n", i);
    }
}

```

To set a conditional breakpoint at the GDB prompt:

```

(gdb) br 8 if i == 8936
Breakpoint 1 at 0x80483f5: file iterations.c, line 8.
(gdb) r

```

With this condition, the program execution will eventually stop with the following output:

```

i = 8931
i = 8932
i = 8933
i = 8934
i = 8935

Breakpoint 1, main () at iterations.c:8
8      fprintf (stdout, "i = %d\n", i);

```

Inspect the breakpoint information (using **info br**) to review the breakpoint status:

```

(gdb) info br
Num      Type           Disp Enb Address      What
1        breakpoint      keep y  0x080483f5  in main at iterations.c:8
          stop only if i == 8936
          breakpoint already hit 1 time

```

#### 4.3.4. Forked Execution

Among the more challenging bugs confronting programmers is where one program (the *parent*) makes an independent copy of itself (a *fork*). That fork then creates a *child* process which, in turn, fails. Debugging the parent process may or may not be useful. Often the only way to get to the bug may be by debugging the child process, but this is not always possible.

The **set follow-fork-mode** feature is used to overcome this barrier allowing programmers to follow a child process instead of the parent process.

##### **set follow-fork-mode parent**

The original process is debugged after a fork. The child process runs unimpeded. This is the default.

##### **set follow-fork-mode child**

The new process is debugged after a fork. The parent process runs unimpeded.

##### **show follow-fork-mode**

Display the current debugger response to a fork call.

Use the **set detach-on-fork** command to debug both the parent and the child processes after a fork, or retain debugger control over them both.

### set detach-on-fork on

The child process (or parent process, depending on the value of **follow-fork-mode**) will be detached and allowed to run independently. This is the default.

### set detach-on-fork off

Both processes will be held under the control of GDB. One process (child or parent, depending on the value of **follow-fork-mode**) is debugged as usual, while the other is suspended.

### show detach-on-fork

Show whether **detach-on-fork** mode is on or off.

Consider the following program:

#### fork.c

```
#include <unistd.h>

int main()
{
    pid_t pid;
    const char *name;

    pid = fork();
    if (pid == 0)
    {
        name = "I am the child";
    }
    else
    {
        name = "I am the parent";
    }
    return 0;
}
```

This program, compiled with the command **gcc -g fork.c -o fork -lpthread** and examined under GDB will show:

```
gdb ./fork
[...]
(gdb) break main
Breakpoint 1 at 0x4005dc: file fork.c, line 8.
(gdb) run
[...]
Breakpoint 1, main () at fork.c:8
8   pid = fork();
(gdb) next
Detaching after fork from child process 3840.
9   if (pid == 0)
(gdb) next
15      name = "I am the parent";
```

```
(gdb) next
17  return 0;
(gdb) print name
$1 = 0x400717 "I am the parent"
```

GDB followed the parent process and allowed the child process (process 3840) to continue execution.

The following is the same test using **set follow-fork-mode child**.

```
(gdb) set follow-fork-mode child
(gdb) break main
Breakpoint 1 at 0x4005dc: file fork.c, line 8.
(gdb) run
[...]
Breakpoint 1, main () at fork.c:8
8  pid = fork();
(gdb) next
[New process 3875]
[Thread debugging using libthread_db enabled]
[Switching to Thread 0x7ffff7fd5720 (LWP 3875)]
9  if (pid == 0)
(gdb) next
11  name = "I am the child";
(gdb) next
17  return 0;
(gdb) print name
$2 = 0x400708 "I am the child"
(gdb)
```

GDB switched to the child process here.

This can be permanent by adding the setting to the appropriate **.gdbinit**.

For example, if **set follow-fork-mode ask** is added to **~/.gdbinit**, then ask mode becomes the default mode.

### 4.3.5. Debugging Individual Threads

GDB has the ability to debug individual threads, and to manipulate and examine them independently. This functionality is not enabled by default. To do so use **set non-stop on** and **set target-async on**. These can be added to **.gdbinit**. Once that functionality is turned on, GDB is ready to conduct thread debugging.

For example, the following program creates two threads. These two threads, along with the original thread executing main makes a total of three threads.

#### three-threads.c

```
#include <stdio.h>
#include <pthread.h>
#include <unistd.h>

pthread_t thread;

void* thread3 (void* d)
```

```
{
    int count3 = 0;

    while(count3 < 1000){
        sleep(10);
        printf("Thread 3: %d\n", count3++);
    }
    return NULL;
}

void* thread2 (void* d)
{
    int count2 = 0;

    while(count2 < 1000){
        printf("Thread 2: %d\n", count2++);
    }
    return NULL;
}

int main (){

    pthread_create (&thread, NULL, thread2, NULL);
    pthread_create (&thread, NULL, thread3, NULL);

    //Thread 1
    int count1 = 0;

    while(count1 < 1000){
        printf("Thread 1: %d\n", count1++);
    }

    pthread_join(thread, NULL);
    return 0;
}
```

Compile this program in order to examine it under GDB.

```
gcc -g three-threads.c -o three-threads -lpthread
gdb ./three-threads
```

First set breakpoints on all thread functions; thread1, thread2, and main.

```
(gdb) break thread3
Breakpoint 1 at 0x4006c0: file three-threads.c, line 9.
(gdb) break thread2
Breakpoint 2 at 0x40070c: file three-threads.c, line 20.
(gdb) break main
Breakpoint 3 at 0x40074a: file three-threads.c, line 30.
```

Then run the program.

```
(gdb) run
[...]
Breakpoint 3, main () at three-threads.c:30
```

```

30 pthread_create (&thread, NULL, thread2, NULL);
[...]
(gdb) info threads
* 1 Thread 0x7ffff7fd5720 (LWP 4620) main () at three-threads.c:30
(gdb)

```

Note that the command **info threads** provides a summary of the program's threads and some details about their current state. In this case there is only one thread that has been created so far.

Continue execution some more.

```

(gdb) next
[New Thread 0x7ffff7fd3710 (LWP 4687)]
31 pthread_create (&thread, NULL, thread3, NULL);
(gdb)
Breakpoint 2, thread2 (d=0x0) at three-threads.c:20
20 int count2 = 0;
next
[New Thread 0x7ffff75d2710 (LWP 4688)]
34 int count1 = 0;
(gdb)
Breakpoint 1, thread3 (d=0x0) at three-threads.c:9
9 int count3 = 0;
info threads
 3 Thread 0x7ffff75d2710 (LWP 4688) thread3 (d=0x0) at three-threads.c:9
 2 Thread 0x7ffff7fd3710 (LWP 4687) thread2 (d=0x0) at three-
threads.c:20
* 1 Thread 0x7ffff7fd5720 (LWP 4620) main () at three-threads.c:34

```

Here, two more threads are created. The star indicates the thread currently under focus. Also, the newly created threads have hit the breakpoint set for them in their initialization functions. Namely, `thread2()` and `thread3()`.

To begin real thread debugging, use the **thread <thread number>** command to switch the focus to another thread.

```

(gdb) thread 2
[Switching to thread 2 (Thread 0x7ffff7fd3710 (LWP 4687))]#0 thread2
(d=0x0)
  at three-threads.c:20
20 int count2 = 0;
(gdb) list
15 return NULL;
16 }
17
18 void* thread2 (void* d)
19 {
20 int count2 = 0;
21
22 while(count2 < 1000){
23 printf("Thread 2: %d\n", count2++);
24 }

```

Thread 2 stopped at line 20 in its function thread2().

```
(gdb) next
22  while(count2 < 1000){
(gdb) print count2
$1 = 0
(gdb) next
23  printf("Thread 2: %d\n", count2++);
(gdb) next
Thread 2: 0
22  while(count2 < 1000){
(gdb) next
23  printf("Thread 2: %d\n", count2++);
(gdb) print count2
$2 = 1
(gdb) info threads
  3 Thread 0x7ffff75d2710 (LWP 4688)  thread3 (d=0x0) at three-threads.c:9
 * 2 Thread 0x7ffff7fd3710 (LWP 4687)  thread2 (d=0x0) at three-
threads.c:23
  1 Thread 0x7ffff7fd5720 (LWP 4620)  main () at three-threads.c:34
(gdb)
```

Above, a few lines of thread2 printed the counter count2 and left thread 2 at line 23 as is seen by the output of 'info threads'.

Now thread3.

```
(gdb) thread 3
[Switching to thread 3 (Thread 0x7ffff75d2710 (LWP 4688))]#0  thread3
(d=0x0)
  at three-threads.c:9
 9  int count3 = 0;
(gdb) list
4
5  pthread_t thread;
6
7  void* thread3 (void* d)
8  {
9  int count3 = 0;
10
11 while(count3 < 1000){
12 sleep(10);
13 printf("Thread 3: %d\n", count3++);
(gdb)
```

Thread three is a little different in that it has a sleep statement and executes slowly. Think of it as a representation of an uninteresting IO thread. Because this thread is uninteresting, continue its execution uninterrupted, using the **continue**.

```
(gdb) continue &
(gdb) Thread 3: 0
Thread 3: 1
Thread 3: 2
Thread 3: 3
```



Take note of the `&` at the end of the **continue**. This allows the GDB prompt to return so other commands can be executed. Using the **interrupt**, execution can be stopped should thread 3 become interesting again.

```
(gdb) interrupt
[Thread 0x7ffff75d2710 (LWP 4688)] #3 stopped.
0x000000343f4a6a6d in nanosleep () at
../sysdeps/unix/syscall-template.S:82
82 T_PSEUDO (SYSCALL_SYMBOL, SYSCALL_NAME, SYSCALL_NARGS)
```

It is also possible to go back to the original main thread and examine it some more.

```
(gdb) thread 1
[Switching to thread 1 (Thread 0x7ffff7fd5720 (LWP 4620))]#0 main ()
    at three-threads.c:34
34 int count1 = 0;
(gdb) next
36 while(count1 < 1000){
(gdb) next
37 printf("Thread 1: %d\n", count1++);
(gdb) next
Thread 1: 0
36 while(count1 < 1000){
(gdb) next
37 printf("Thread 1: %d\n", count1++);
(gdb) next
Thread 1: 1
36 while(count1 < 1000){
(gdb) next
37 printf("Thread 1: %d\n", count1++);
(gdb) next
Thread 1: 2
36 while(count1 < 1000){
(gdb) print count1
$3 = 3
(gdb) info threads
  3 Thread 0x7ffff75d2710 (LWP 4688) 0x000000343f4a6a6d in nanosleep ()
    at ../sysdeps/unix/syscall-template.S:82
  2 Thread 0x7ffff7fd3710 (LWP 4687) thread2 (d=0x0) at three-
threads.c:23
* 1 Thread 0x7ffff7fd5720 (LWP 4620) main () at three-threads.c:36
(gdb)
```

As can be seen from the output of `info threads`, the other threads are where they were left, unaffected by the debugging of thread 1.

### 4.3.6. Alternative User Interfaces for GDB

GDB uses the command line as its default interface. However, it also has an API called *machine interface* (MI). MI allows IDE developers to create other user interfaces to GDB.

Some examples of these interfaces are:

#### Eclipse (CDT)

A graphical debugger interface integrated with the Eclipse development environment. More information can be found at the [Eclipse website](#).

### Nemiver

A graphical debugger interface which is well suited to the GNOME Desktop Environment. More information can be found at the [Nemiver website](#)

### Emacs

A GDB interface which is integrated with the emacs. More information can be found at the [Emacs website](#)

## 4.4. VARIABLE TRACKING AT ASSIGNMENTS

*Variable Tracking at Assignments* (VTA) is a new infrastructure included in GCC used to improve variable tracking during optimizations. This allows GCC to produce more precise, meaningful, and useful debugging information for GDB, SystemTap, and other debugging tools.

When GCC compiles code with optimizations enabled, variables are renamed, moved around, or even removed altogether. As such, optimized compiling can cause a debugger to report that some variables have been <optimized out>. With VTA enabled, optimized code is internally annotated to ensure that optimization passes to transparently keep track of each variable's value, regardless of whether the variable is moved or removed. The effect of this is more parameter and variable values available, even for the optimized (`gcc -O2 -g` built) code. It also displays the <optimized out> message less.

VTA's benefits are more pronounced when debugging applications with inlined functions. Without VTA, optimization could completely remove some arguments of an inlined function, preventing the debugger from inspecting its value. With VTA, optimization will still happen, and appropriate debugging information will be generated for any missing arguments.

VTA is enabled by default when compiling code with optimizations and debugging information enabled (that is, `gcc -O -g` or, more commonly, `gcc -O2 -g`). To disable VTA during such builds, add the `-fno-var-tracking-assignments`. In addition, the VTA infrastructure includes the new `gcc` option `-fcompare-debug`. This option tests code compiled by GCC with debug information and without debug information: the test passes if the two binaries are identical. This test ensures that executable code is not affected by any debugging options, which further ensures that there are no hidden bugs in the debug code. Note that `-fcompare-debug` adds significant cost in compilation time. See `man gcc` for details about this option.

For more information about the infrastructure and development of VTA, see *A Plan to Fix Local Variable Debug Information in GCC*, available at the following link:

[http://gcc.gnu.org/wiki/Var\\_Tracking\\_Assignments](http://gcc.gnu.org/wiki/Var_Tracking_Assignments)

A slide deck version of this whitepaper is also available at <http://people.redhat.com/aoliva/papers/vta/slides.pdf>.

## 4.5. PYTHON PRETTY-PRINTERS

The GDB command `print` outputs comprehensive debugging information for a target application. GDB aims to provide as much debugging data as it can to users; however, this means that for highly complex programs the amount of data can become very cryptic.

In addition, GDB does not provide any tools that help decipher GDB **print** output. GDB does not even empower users to easily create tools that can help decipher program data. This makes the practice of reading and understanding debugging data quite arcane, particularly for large, complex projects.

For most developers, the only way to customize GDB **print** output (and make it more meaningful) is to revise and recompile GDB. However, very few developers can actually do this. Further, this practice will not scale well, particularly if the developer must also debug other programs that are heterogeneous and contain equally complex debugging data.

To address this, the Red Hat Enterprise Linux 6 version of GDB is now compatible with Python *pretty-printers*. This allows the retrieval of more meaningful debugging data by leaving the introspection, printing, and formatting logic to a *third-party* Python script.

Compatibility with Python pretty-printers gives you the chance to truly customize GDB output as you see fit. This makes GDB a more viable debugging solution to a wider range of projects, since you now have the flexibility to *adapt* GDB output as required, and with greater ease. Further, developers with intimate knowledge of a project and a specific programming language are best qualified in deciding what kind of output is meaningful, allowing them to improve the usefulness of that output.

The Python pretty-printers implementation allows users to automatically inspect, format, and print program data according to specification. These specifications are written as rules implemented via Python scripts. This offers the following benefits:

### Safe

To pass program data to a set of registered Python pretty-printers, the GDB development team added *hooks* to the GDB printing code. These hooks were implemented with safety in mind: the built-in GDB printing code is still intact, allowing it to serve as a default fallback printing logic. As such, if no specialized printers are available, GDB will still print debugging data the way it always did. This ensures that GDB is backwards-compatible; users who do not require pretty-printers can still continue using GDB.

### Highly Customizable

This new "Python-scripted" approach allows users to distill as much knowledge as required into specific printers. As such, a project can have an entire library of printer scripts that parses program data in a unique manner specific to its user's requirements. There is no limit to the number of printers a user can build for a specific project; what's more, being able to customize debugging data script by script offers users an easier way to re-use and re-purpose printer scripts — or even a whole library of them.

### Easy to Learn

The best part about this approach is its lower barrier to entry. Python scripting is comparatively easy to learn and has a large library of free documentation available online. In addition, most programmers already have basic to intermediate experience in Python scripting, or in scripting in general.

Here is a small example of a pretty printer. Consider the following C++ program:

#### fruit.cc

```
enum Fruits {Orange, Apple, Banana};

class Fruit
{
    int fruit;

public:
    Fruit (int f)
    {
```

```
        fruit = f;
    }
};

int main()
{
    Fruit myFruit(Apple);
    return 0;           // line 17
}
```

This is compiled with the command `g++ -g fruit.cc -o fruit`. Now, examine this program with GDB.

```
gdb ./fruit
[...]
(gdb) break 17
Breakpoint 1 at 0x40056d: file fruit.cc, line 17.
(gdb) run

Breakpoint 1, main () at fruit.cc:17
17  return 0;           // line 17
(gdb) print myFruit
$1 = {fruit = 1}
```

The output of `{fruit = 1}` is correct because that is the internal representation of 'fruit' in the data structure 'Fruit'. However, this is not easily read by humans as it is difficult to tell which fruit the integer 1 represents.

To solve this problem, write the following pretty printer:

```
fruit.py

class FruitPrinter:
    def __init__(self, val):
        self.val = val

    def to_string (self):
        fruit = self.val['fruit']

        if (fruit == 0):
            name = "Orange"
        elif (fruit == 1):
            name = "Apple"
        elif (fruit == 2):
            name = "Banana"
        else:
            name = "unknown"
        return "Our fruit is " + name

def lookup_type (val):
    if str(val.type) == 'Fruit':
        return FruitPrinter(val)
    return None

gdb.pretty_printers.append (lookup_type)
```

Examine this printer from the bottom up.

The line `gdb.pretty_printers.append (lookup_type)` adds the function `lookup_type` to GDB's list of printer lookup functions.

The function `lookup_type` is responsible for examining the type of object to be printed, and returning an appropriate pretty printer. The object is passed by GDB in the parameter `val`. `val.type` is an attribute which represents the type of the pretty printer.

`FruitPrinter` is where the actual work is done. More specifically in the `to_string` function of that Class. In this function, the integer `fruit` is retrieved using the python dictionary syntax `self.val['fruit']`. Then the name is determined using that value. The string returned by this function is the string that will be printed to the user.

After creating `fruit.py`, it must then be loaded into GDB with the following command:

```
(gdb) python execfile("fruit.py")
```

The *GDB and Python Pretty-Printers* whitepaper provides more details on this feature. This whitepaper also includes details and examples on how to write your own Python pretty-printer as well as how to import it into GDB. See the following link for more information:

<http://sourceware.org/gdb/onlinedocs/gdb/Pretty-Printing.html>

## CHAPTER 5. PROFILING

Developers profile programs to focus attention on the areas of the program that have the largest impact on performance. The types of data collected include what section of the program consumes the most processor time, and where memory is allocated. Profiling collects data from the actual program execution. Thus, the quality of the data collect is influenced by the actual tasks being performed by the program. The tasks performed during profiling should be representative of actual use; this ensures that problems arising from realistic use of the program are addressed during development.

Red Hat Enterprise Linux 6 includes a number of different tools (Valgrind, OProfile, **perf**, and SystemTap) to collect profiling data. Each tool is suitable for performing specific types of profile runs, as described in the following sections.

### 5.1. VALGRIND

**Valgrind** is an instrumentation framework for building dynamic analysis tools that can be used to profile applications in detail. The default installation already provides five standard tools. **Valgrind** tools are generally used to investigate memory management and threading problems. The **Valgrind** suite also includes tools that allow the building of new profiling tools as required.

**Valgrind** provides instrumentation for user-space binaries to check for errors, such as the use of uninitialized memory, improper allocation/freeing of memory, and improper arguments for systemcalls. Its profiling tools can be used by normal users on most binaries; however, compared to other profilers, **Valgrind** profile runs are significantly slower. To profile a binary, **Valgrind** rewrites its executable and instruments the rewritten binary. **Valgrind**'s tools are most useful for looking for memory-related issues in user-space programs; it is not suitable for debugging time-specific issues or kernel-space instrumentation and debugging.

**Valgrind** reports are most useful and accurate when *debuginfo* packages are installed for the programs or libraries under investigation. See [Section 4.2, “Installing Debuginfo Packages”](#).

#### 5.1.1. Valgrind Tools

The **Valgrind** suite is composed of the following tools:

##### memcheck

This tool detects memory management problems in programs by checking all reads from and writes to memory and intercepting all system calls to **malloc**, **new**, **free**, and **delete**. **memcheck** is perhaps the most used **Valgrind** tool, as memory management problems can be difficult to detect using other means. Such problems often remain undetected for long periods, eventually causing crashes that are difficult to diagnose.

##### cachegrind

**cachegrind** is a cache profiler that accurately pinpoints sources of cache misses in code by performing a detailed simulation of the I1, D1 and L2 caches in the CPU. It shows the number of cache misses, memory references, and instructions accruing to each line of source code; **cachegrind** also provides per-function, per-module, and whole-program summaries, and can even show counts for each individual machine instructions.

##### callgrind

Like **cachegrind**, **callgrind** can model cache behavior. However, the main purpose of **callgrind** is to record callgraphs data for the executed code.

## massif

**massif** is a heap profiler; it measures how much heap memory a program uses, providing information on heap blocks, heap administration overheads, and stack sizes. Heap profilers are useful in finding ways to reduce heap memory usage. On systems that use virtual memory, programs with optimized heap memory usage are less likely to run out of memory, and may be faster as they require less paging.

## helgrind

In programs that use the POSIX pthreads threading primitives, **helgrind** detects synchronization errors. Such errors are:

- Misuses of the POSIX pthreads API
- Potential deadlocks arising from lock ordering problems
- Data races (that is, accessing memory without adequate locking)

**Valgrind** also allows you to develop your own profiling tools. In line with this, **Valgrind** includes the **lackey** tool, which is a sample that can be used as a template for generating your own tools.

### 5.1.2. Using Valgrind

The `valgrind` package and its dependencies install all the necessary tools for performing a **Valgrind** profile run. To profile a program with **Valgrind**, use:

```
~]$ valgrind --tool=toolname program
```

See [Section 5.1.1, “Valgrind Tools”](#) for a list of arguments for *toolname*. In addition to the suite of **Valgrind** tools, **none** is also a valid argument for *toolname*; this argument allows you to run a program under **Valgrind** without performing any profiling. This is useful for debugging or benchmarking **Valgrind** itself.

You can also instruct **Valgrind** to send all of its information to a specific file. To do so, use the option `--log-file=filename`. For example, to check the memory usage of the executable file `hello` and send profile information to `output`, use:

```
~]$ valgrind --tool=memcheck --log-file=output hello
```

See [Section 5.1.3, “Additional information”](#) for more information on **Valgrind**, along with other available documentation on the **Valgrind** suite of tools.

### 5.1.3. Additional information

For more extensive information on **Valgrind**, see `man valgrind`. Red Hat Enterprise Linux also provides a comprehensive *Valgrind Documentation* book available as PDF and HTML in:

- `/usr/share/doc/valgrind-version/valgrind_manual.pdf`
- `/usr/share/doc/valgrind-version/html/index.html`

## 5.2. OPROFILE

**OProfile** is a low overhead, system-wide performance monitoring tool provided by the `oprofile` package. It uses the performance monitoring hardware on the processor to retrieve information about the kernel and executables on the system, such as when memory is referenced, the number of second-level cache requests, and the number of hardware interrupts received. **OProfile** is also able to profile applications that run in a Java Virtual Machine (JVM).

The following is a selection of the tools provided by **OProfile**. Note that the legacy `opcontrol` tool and the new `operf` tool are mutually exclusive.

### **ophelp**

Displays available events for the system's processor along with a brief description of each.

### **operf**

Intended to replace `opcontrol`. The `operf` tool uses the Linux Performance Events subsystem, allowing you to target your profiling more precisely, as a single process or system-wide, and allowing **OProfile** to co-exist better with other tools using the performance monitoring hardware on your system. Unlike `opcontrol`, no initial setup is required, and it can be used without the root privileges unless the `--system-wide` option is in use.

### **opimport**

Converts sample database files from a foreign binary format to the native format for the system. Only use this option when analyzing a sample database from a different architecture.

### **opannotate**

Creates an annotated source for an executable if the application was compiled with debugging symbols.

### **opreport**

Retrieves profile data.

### **opcontrol**

This tool is used to start and stop the **OProfile** daemon (`oprofiled`) and configure a profile session.

### **oprofiled**

Runs as a daemon to periodically write sample data to disk.

Legacy mode (`opcontrol`, `oprofiled`, and post-processing tools) remains available, but it is no longer the recommended profiling method. For a detailed description of the legacy mode, see the [Configuring OProfile Using Legacy Mode](#) chapter in the *System Administrator's Guide*.

## 5.2.1. Using OProfile

`operf` is the recommended tool for collecting profiling data. The tool does not require any initial configuration, and all options are passed to it on the command line. Unlike the legacy `opcontrol` tool, `operf` can run without `root` privileges. See the [Using operf](#) chapter in the *System Administrator's Guide* for detailed instructions on how to use the `operf` tool.

### Example 5.1. Using operf to Profile a Java Program



In the following example, the **operf** tool is used to collect profiling data from a Java (JIT) program, and the **opreport** tool is then used to output per-symbol data.

1. Install the demonstration Java program used in this example. It is a part of the java-1.8.0-openjdk-demo package, which is included in the *Optional* channel. See [Enabling Supplementary and Optional Repositories](#) for instructions on how to use the *Optional* channel. When the *Optional* channel is enabled, install the package:

```
~]# yum install java-1.8.0-openjdk-demo
```

2. Install the oprofile-jit package for **OProfile** to be able to collect profiling data from Java programs:

```
~]# yum install oprofile-jit
```

3. Create a directory for **OProfile** data:

```
~]$ mkdir ~/oprofile_data
```

4. Change into the directory with the demonstration program:

```
~]$ cd /usr/lib/jvm/java-1.8.0-  
openjdk/demo/applets/MoleculeViewer/
```

5. Start the profiling:

```
~]$ operf -d ~/oprofile_data appletviewer \  
-J"-agentpath:/usr/lib64/oprofile/libjvmti_oprofile.so"  
example2.html
```

6. Change into the home directory and analyze the collected data:

```
~]$ cd
```

```
~]$ opreport --symbols --threshold 0.5
```

A sample output may look like the following:

```
$ opreport --symbols --threshold 0.5  
Using /home/rkratky/oprofile_data/samples/ for samples directory.  
  
WARNING! Some of the events were throttled. Throttling occurs when  
the initial sample rate is too high, causing an excessive number  
of  
interrupts. Decrease the sampling frequency. Check the directory  
/home/rkratky/oprofile_data/samples/current/stats/throttled  
for the throttled event names.  
  
warning: /dm_crypt could not be found.  
warning: /e1000e could not be found.  
warning: /kvm could not be found.  
CPU: Intel Ivy Bridge microarchitecture, speed 3600 MHz
```

```

(estimated)
Counted CPU_CLK_UNHALTED events (Clock cycles when not halted)
with a unit mask of 0x00 (No unit mask) count 100000
samples %      image name                symbol name
14270   57.1257  libjvm.so                               /usr/lib/jvm/java-
1.8.0-openjdk-1.8.0.51-
1.b16.e17_1.x86_64/jre/lib/amd64/server/libjvm.so
3537    14.1593  23719.jo                                Interpreter
690     2.7622   libc-2.17.so                            fgetc
581     2.3259   libX11.so.6.3.0                          /usr/lib64/libX11.so.6.3.0
364     1.4572   libpthread-2.17.so                       pthread_getspecific
130     0.5204   libfreetype.so.6.10.0                    /usr/lib64/libfreetype.so.6.10.0
128     0.5124   libc-2.17.so                             __memset_sse2

```

## 5.2.2. OProfile in Red Hat Enterprise Linux 7

**OProfile 0.9.9**, which is included in Red Hat Enterprise Linux 7, contains a number of improvements over previous versions. The new version also supports the utilization of the *Linux Performance Events* subsystem using the new **opperf** command.

### 5.2.2.1. New Features

A new **opperf** program is now available that allows non-root users to profile single processes. This can also be used for system-wide profiling, but in this case, root authority is required.

**OProfile 0.9.9** now supports the following:

- IBM POWER8 processors
- Intel Haswell processors
- IBM zEnterprise EC12 (zEC12) processor
- AMD Generic Performance Events
- IBM Power ISA 2.07 Architected Events

### 5.2.2.2. Known Problems and Limitations

**OProfile 0.9.9** has a few known problems and limitations. These are:

- AMD Instruction Based Sampling (IBS) is not currently supported with the new **opperf** program. Use the legacy **opcontrol** commands for IBS profiling.
- The type of the sample header **mtime** field has changed to u64, which makes it impossible to process sample data acquired using previous versions of **OProfile**.
- **opcontrol** fails to allocate the hardware performance counters it needs if the *NMI watchdog* is enabled. The NMI watchdog, which monitors system interrupts, uses the **perf** tool, which reserves all performance counters.

### 5.2.3. OProfile Documentation

For more extensive information on **OProfile**, see the `oprofile(1)` manual page. Red Hat Enterprise Linux also provides two comprehensive guides to **OProfile** in `file:///usr/share/doc/oprofile-version/`:

#### OProfile Manual

A comprehensive manual with detailed instructions on the setup and use of **OProfile** is found at `file:///usr/share/doc/oprofile-version/oprofile.html`

#### OProfile Internals

Documentation on the internal workings of **OProfile**, useful for programmers interested in contributing to the **OProfile** upstream, can be found at `file:///usr/share/doc/oprofile-version/internals.html`

## 5.3. SYSTEMTAP

SystemTap is a useful instrumentation platform for probing running processes and kernel activity on the Linux system. To execute a probe:

1. Write *SystemTap scripts* that specify which system events (for example, virtual file system reads, packet transmissions) should trigger specified actions (for example, print, parse, or otherwise manipulate data).
2. SystemTap translates the script into a C program, which it compiles into a kernel module.
3. SystemTap loads the kernel module to perform the actual probe.

SystemTap scripts are useful for monitoring system operation and diagnosing system issues with minimal intrusion into the normal operation of the system. You can quickly instrument running system test hypotheses without having to recompile and re-install instrumented code. To compile a SystemTap script that probes *kernel-space*, SystemTap uses information from three different *kernel information packages*:

- `kernel-variant-devel-version`
- `kernel-variant-debuginfo-version`
- `kernel-debuginfo-common-arch-version`



#### NOTE

The kernel information package in Red Hat Enterprise Linux 6 is now named `kernel-debuginfo-common-arch-version`. It was originally `kernel-debuginfo-common-version` in Red Hat Enterprise Linux 5.

These kernel information packages must match the kernel to be probed. In addition, to compile SystemTap scripts for multiple kernels, the kernel information packages of each kernel must also be installed.

### 5.3.1. Additional Information

For more detailed information about SystemTap, see the following Red Hat documentation:

- [SystemTap Beginner's Guide](#)
- [SystemTap Tapset Reference](#)

## 5.4. PERFORMANCE COUNTERS FOR LINUX (PCL) TOOLS AND PERF

*Performance Counters for Linux* (PCL) is a new kernel-based subsystem that provides a framework for collecting and analyzing performance data. These events will vary based on the performance monitoring hardware and the software configuration of the system. Red Hat Enterprise Linux 6 includes this kernel subsystem to collect data and the user-space tool **perf** to analyze the collected performance data.

The PCL subsystem can be used to measure hardware events, including retired instructions and processor clock cycles. It can also measure software events, including major page faults and context switches. For example, PCL counters can compute the *Instructions Per Clock* (IPC) from a process's counts of instructions retired and processor clock cycles. A low IPC ratio indicates the code makes poor use of the CPU. Other hardware events can also be used to diagnose poor CPU performance.

Performance counters can also be configured to record samples. The relative frequency of samples can be used to identify which regions of code have the greatest impact on performance.

### 5.4.1. Perf Tool Commands

Useful **perf** commands include the following:

#### **perf stat**

This **perf** command provides overall statistics for common performance events, including instructions executed and clock cycles consumed. Options allow selection of events other than the default measurement events.

#### **perf record**

This **perf** command records performance data into a file which can be later analyzed using **perf report**.

#### **perf report**

This **perf** command reads the performance data from a file and analyzes the recorded data.

#### **perf list**

This **perf** command lists the events available on a particular machine. These events will vary based on the performance monitoring hardware and the software configuration of the system.

Use **perf help** to obtain a complete list of **perf** commands. To retrieve **man** page information on each **perf** command, use **perf help command**.

### 5.4.2. Using Perf

Using the basic PCL infrastructure for collecting statistics or samples of program execution is relatively straightforward. This section provides simple examples of overall statistics and sampling.

To collect statistics on **make** and its children, use the following command:

```
# perf stat -- make all
```

The **perf** command collects a number of different hardware and software counters. It then prints the following information:

```
Performance counter stats for 'make all':

 244011.782059 task-clock-msecs      #      0.925 CPUs
      53328 context-switches        #      0.000 M/sec
      515 CPU-migrations             #      0.000 M/sec
 1843121 page-faults                #      0.008 M/sec
 789702529782 cycles                 #    3236.330 M/sec
1050912611378 instructions           #      1.331 IPC
 275538938708 branches              #    1129.203 M/sec
 2888756216 branch-misses           #      1.048 %
 4343060367 cache-references        #     17.799 M/sec
 428257037 cache-misses             #      1.755 M/sec

263.779192511 seconds time elapsed
```

The **perf** tool can also record samples. For example, to record data on the **make** command and its children, use:

```
# perf record -- make all
```

This prints out the file in which the samples are stored, along with the number of samples collected:

```
[ perf record: Woken up 42 times to write data ]
[ perf record: Captured and wrote 9.753 MB perf.data (~426109 samples) ]
```

As of Red Hat Enterprise Linux 6.4, a new functionality to the **{}** group syntax has been added that allows the creation of event groups based on the way they are specified on the command line.

The current **--group** or **-g** options remain the same; if it is specified for **record**, **stat**, or **top** command, all the specified events become members of a single group with the first event as a group leader.

The new **{}** group syntax allows the creation of a group like:

```
# perf record -e '{cycles, faults}' ls
```

The above results in a single event group containing *cycles* and *faults* events, with the *cycles* event as the group leader.

All groups are created with regards to threads and CPUs. As such, recording an event group within two threads on a server with four CPUs will create eight separate groups.

It is possible to use a standard event modifier for a group. This spans over all events in the group and updates each event modifier settings.

```
# perf record -r '{faults:k, cache-references}:p'
```

The above command results in the **:kp** modifier being used for *faults*, and the **:p** modifier being used for the *cache-references* event.

## Performance Counters for Linux (PCL) Tools conflict with OProfile

Both OProfile and Performance Counters for Linux (PCL) use the same hardware Performance Monitoring Unit (PMU). If OProfile is currently running while attempting to use the PCL **perf** command, an error message like the following occurs when starting OProfile:

```
Error: open_counter returned with 16 (Device or resource busy). /bin/dmesg
may provide additional information.
```

```
Fatal: Not all events could be opened.
```

To use the **perf** command, first shut down OProfile:

```
# opcontrol --deinit
```

You can then analyze **perf.data** to determine the relative frequency of samples. The report output includes the command, object, and function for the samples. Use **perf report** to output an analysis of **perf.data**. For example, the following command produces a report of the executable that consumes the most time:

```
# perf report --sort=comm
```

The resulting output:

```
# Samples: 1083783860000
#
# Overhead      Command
# .....
#
 48.19%        xsltproc
 44.48%        pdfxmltex
  6.01%         make
  0.95%         perl
  0.17%        kernel-doc
  0.05%         xmllint
  0.05%         cc1
  0.03%         cp
  0.01%         xmlto
  0.01%         sh
  0.01%        docproc
  0.01%         ld
  0.01%         gcc
  0.00%         rm
  0.00%         sed
  0.00%        git-diff-files
  0.00%         bash
  0.00%        git-diff-index
```

The column on the left shows the relative frequency of the samples. This output shows that **make** spends most of this time in **xsltproc** and the **pdfxmltex**. To reduce the time for the **make** to complete, focus on **xsltproc** and **pdfxmltex**. To list the functions executed by **xsltproc**, run:

```
# perf report -n --comm=xsltproc
```

This generates:

```
-
```

```

comm: xsltproc
# Samples: 472520675377
#
# Overhead  Samples                Shared Object  Symbol
# .....
#
 45.54%215179861044  libxml2.so.2.7.6      [.]
xmlXPathCmpNodesExt
 11.63%54959620202  libxml2.so.2.7.6      [.]
xmlXPathNodeSetAdd__internal_alias
  8.60%40634845107  libxml2.so.2.7.6      [.]
xmlXPathCompOpEval
  4.63%21864091080  libxml2.so.2.7.6      [.]
xmlXPathReleaseObject
  2.73%12919672281  libxml2.so.2.7.6      [.]
xmlXPathNodeSetSort__internal_alias
  2.60%12271959697  libxml2.so.2.7.6      [.] valuePop
  2.41%11379910918  libxml2.so.2.7.6      [.]
xmlXPathIsNaN__internal_alias
  2.19%10340901937  libxml2.so.2.7.6      [.]
valuePush__internal_alias

```

## 5.5. FTRACE

The **ftrace** framework provides users with several tracing capabilities, accessible through an interface much simpler than SystemTap's. This framework uses a set of virtual files in the **debugfs** file system; these files enable specific tracers. The **ftrace** function tracer outputs each function called in the kernel in real time; other tracers within the **ftrace** framework can also be used to analyze wakeup latency, task switches, kernel events, and the like.

You can also add new tracers for **ftrace**, making it a flexible solution for analyzing kernel events. The **ftrace** framework is useful for debugging or analyzing latencies and performance issues that take place outside of user-space. Unlike other profilers documented in this guide, **ftrace** is a built-in feature of the kernel.

### 5.5.1. Using ftrace

The Red Hat Enterprise Linux 6 kernels have been configured with the **CONFIG\_FTRACE=y** option. This option provides the interfaces required by **ftrace**. To use **ftrace**, mount the **debugfs** file system as follows:

```
mount -t debugfs nodev /sys/kernel/debug
```

All the **ftrace** utilities are located in **/sys/kernel/debug/tracing/**. View the **/sys/kernel/debug/tracing/available\_tracers** file to find out what tracers are available for your kernel:

```
cat /sys/kernel/debug/tracing/available_tracers
```

```
power wakeup irqsoff function sysprof sched_switch initcall nop
```

To use a specific tracer, write it to `/sys/kernel/debug/tracing/current_tracer`. For example, **wakeup** traces and records the maximum time it takes for the highest-priority task to be scheduled after the task wakes up. To use it:

```
echo wakeup > /sys/kernel/debug/tracing/current_tracer
```

To start or stop tracing, write to `/sys/kernel/debug/tracing/tracing_on`, as in:

```
echo 1 > /sys/kernel/debug/tracing/tracing_on (enables tracing)
```

```
echo 0 > /sys/kernel/debug/tracing/tracing_on (disables tracing)
```

The results of the trace can be viewed from the following files:

**`/sys/kernel/debug/tracing/trace`**

This file contains human-readable trace output.

**`/sys/kernel/debug/tracing/trace_pipe`**

This file contains the same output as `/sys/kernel/debug/tracing/trace`, but is meant to be piped into a command. Unlike `/sys/kernel/debug/tracing/trace`, reading from this file consumes its output.

## 5.5.2. ftrace Documentation

The **ftrace** framework is fully documented in the following files:

- *ftrace - Function Tracer*: **file:///usr/share/doc/kernel-doc-version/Documentation/trace/ftrace.txt**
- *function tracer guts*: **file:///usr/share/doc/kernel-doc-version/Documentation/trace/ftrace-design.txt**



## CHAPTER 6. DOCUMENTATION TOOLS

Red Hat Enterprise Linux 6 offers the **Doxygen** tool for generating documentation from source code and for writing standalone documentation.

### 6.1. DOXYGEN

Doxygen is a documentation tool that creates reference material both online in HTML and offline in Latex. It does this from a set of documented source files which makes it easy to keep the documentation consistent and correct with the source code.

#### 6.1.1. Doxygen Supported Output and Languages

Doxygen has support for output in:

- RTF (MS Word)
- PostScript
- Hyperlinked PDF
- Compressed HTML
- Unix man pages

Doxygen supports the following programming languages:

- C
- C++
- C#
- Objective -C
- IDL
- Java
- VHDL
- PHP
- Python
- Fortran
- D

#### 6.1.2. Getting Started

Doxygen uses a configuration file to determine its settings, therefore it is paramount that this be created correctly. Each project requires its own configuration file. The most painless way to create the configuration file is with the command **doxygen -g *config-file***. This creates a template configuration file that can be easily edited. The variable *config-file* is the name of the configuration file. If

it is committed from the command it is called Doxyfile by default. Another useful option while creating the configuration file is the use of a minus sign (-) as the file name. This is useful for scripting as it will cause Doxygen to attempt to read the configuration file from standard input (**stdin**).

The configuration file consists of a number of variables and tags, similar to a simple Makefile. For example:

```
TAGNAME = VALUE1 VALUE2 . . .
```

For the most part these can be left alone but should it be required to edit them see the [configuration page](#) of the Doxygen documentation website for an extensive explanation of all the tags available. There is also a GUI interface called **doxywizard**. If this is the preferred method of editing then documentation for this function can be found on the [Doxywizard usage page](#) of the Doxygen documentation website.

There are eight tags that are useful to become familiar with.

### **INPUT**

For small projects consisting mainly of C or C++ source and header files it is not required to change anything. However, if the project is large and consists of a source directory or tree, then assign the root directory or directories to the INPUT tag.

### **FILE\_PATTERNS**

File patterns (for example, \*.**cpp** or \*.**h**) can be added to this tag allowing only files that match one of the patterns to be parsed.

### **RECURSIVE**

Setting this to **yes** will allow recursive parsing of a source tree.

### **EXCLUDE and EXCLUDE\_PATTERNS**

These are used to further fine-tune the files that are parsed by adding file patterns to avoid. For example, to omit all **test** directories from a source tree, use **EXCLUDE\_PATTERNS = \*/test/\***.

### **EXTRACT\_ALL**

When this is set to **yes**, doxygen will pretend that everything in the source files is documented to give an idea of how a fully documented project would look. However, warnings regarding undocumented members will not be generated in this mode; set it back to **no** when finished to correct this.

### **SOURCE\_BROWSER and INLINE\_SOURCES**

By setting the **SOURCE\_BROWSER** tag to **yes** doxygen will generate a cross-reference to analyze a piece of software's definition in its source files with the documentation existing about it. These sources can also be included in the documentation by setting **INLINE\_SOURCES** to **yes**.

## **6.1.3. Running Doxygen**

Running **doxygen config-file** creates **html**, **rtf**, **latex**, **xml**, and / or **man** directories in whichever directory doxygen is started in, containing the documentation for the corresponding filetype.

### **HTML OUTPUT**

This documentation can be viewed with a HTML browser that supports cascading style sheets (CSS), as well as DHTML and Javascript for some sections. Point the browser (for example, Mozilla, Safari, Konqueror, or Internet Explorer 6) to the **index.html** in the **html** directory.

## LaTeX OUTPUT

Doxygen writes a **Makefile** into the **latex** directory in order to make it easy to first compile the LaTeX documentation. To do this, use a recent TeX distribution. What is contained in this directory depends on whether the **USE\_PDFLATEX** is set to **no**. Where this is true, typing **make** while in the **latex** directory generates **refman.dvi**. This can then be viewed with **xdvi** or converted to **refman.ps** by typing **make ps**. Note that this requires **dvips**.

There are a number of commands that may be useful. The command **make ps\_2on1** prints two pages on one physical page. It is also possible to convert to a PDF if a ghostscript interpreter is installed by using the command **make pdf**. Another valid command is **make pdf\_2on1**. When doing this set **PDF\_HYPERLINKS** and **USE\_PDFLATEX** tags to **yes** as this will cause **Makefile** will only contain a target to build **refman.pdf** directly.

## RTF OUTPUT

This file is designed to import into Microsoft Word by combining the RTF output into a single file: **refman.rtf**. Some information is encoded using fields but this can be shown by selecting all (**CTRL+A** or Edit -> select all) and then right-click and select the **toggle fields** option from the drop down menu.

## XML OUTPUT

The output into the **xml** directory consists of a number of files, each compound gathered by doxygen, as well as an **index.xml**. An XSLT script, **combine.xslt**, is also created that is used to combine all the XML files into a single file. Along with this, two XML schema files are created, **index.xsd** for the index file, and **compound.xsd** for the compound files, which describe the possible elements, their attributes, and how they are structured.

## MAN PAGE OUTPUT

The documentation from the **man** directory can be viewed with the **man** program after ensuring the **manpath** has the correct man directory in the man path. Be aware that due to limitations with the man page format, information such as diagrams, cross-references and formulas will be lost.


### 6.1.4. Documenting the Sources

There are three main steps to document the sources.

1. First, ensure that **EXTRACT\_ALL** is set to **no** so warnings are correctly generated and documentation is built properly. This allows doxygen to create documentation for documented members, files, classes and namespaces.
2. There are two ways this documentation can be created:

#### A *special* documentation block

This comment block, containing additional marking so Doxygen knows it is part of the documentation, is in either C or C++. It consists of a brief description, or a detailed description. Both of these are optional. What is not optional, however, is the *in body* description. This then links together all the comment blocks found in the body of the method or function.



#### NOTE

While more than one brief or detailed description is allowed, this is not recommended as the order is not specified.

The following will detail the ways in which a comment block can be marked as a detailed description:

- o C-style comment block, starting with two asterisks (\*) in the JavaDoc style.

```
/**
 * ... documentation ...
 */
```

- o C-style comment block using the Qt style, consisting of an exclamation mark (!) instead of an extra asterisk.

```
/*!
 * ... documentation ...
 */
```

- o The beginning asterisks on the documentation lines can be left out in both cases if that is preferred.
- o A blank beginning and end line in C++ also acceptable, with either three forward slashes or two forward slashes and an exclamation mark.

```
///
/// ... documentation
///
```

or

```
//!
//! ... documentation ...
//!
```

- o Alternatively, in order to make the comment blocks more visible a line of asterisks or forward slashes can be used.

```
////////////////////////////////////
/// ... documentation ...
////////////////////////////////////
```

or

```
/******//**
 * ... documentation ...
*****/
```

Note that the two forwards slashes at the end of the normal comment block start a special comment block.

There are three ways to add a brief description to documentation.

- o To add a brief description use `\brief` above one of the comment blocks. This brief section ends at the end of the paragraph and any further paragraphs are the detailed descriptions.

```

    /*! \brief brief documentation.
     *      brief documentation.
     *
     *      detailed documentation.
     */

```

- o By setting **JAVADOC\_AUTOBRIEF** to **yes**, the brief description will only last until the first dot followed by a space or new line. Consequentially limiting the brief description to a single sentence.

```

    /** Brief documentation. Detailed documentation continues *
     * from here.
     */

```

This can also be used with the above mentioned three-slash comment blocks (///).

- o The third option is to use a special C++ style comment, ensuring this does not span more than one line.

```

    /// Brief documentation.
    /** Detailed documentation. */

```

or

```

    //! Brief documentation.
    //! Detailed documentation //!< starts here

```

The blank line in the above example is required to separate the brief description and the detailed description, and **JAVADOC\_AUTOBRIEF** must to be set to **no**.

Examples of how a documented piece of C++ code using the Qt style can be found on the [Doxygen documentation website](#)

It is also possible to have the documentation after members of a file, struct, union, class, or enum. To do this add a < marker in the comment block.\

```

    int var; /*!< detailed description after the member */

```

Or in a Qt style as:

```

    int var; /**< detailed description after the member */

```

or

```

    int var; //!< detailed description after the member
           //!<

```

or

```

    int var; ///< detailed description after the member
           ///<

```

For brief descriptions after a member use:

```
int var; //!< brief description after the member
```

or

```
int var; /// brief description after the member
```

Examples of these and how the HTML is produced can be viewed on the [Doxygen documentation website](#)

## Documentation at other places

While it is preferable to place documentation in front of the code it is documenting, at times it is only possible to put it in a different location, especially if a file is to be documented; after all it is impossible to place the documentation in front of a file. This is best avoided unless it is absolutely necessary as it can lead to some duplication of information.

To do this it is important to have a structural command inside the documentation block. Structural commands start with a backslash (\) or an at-sign (@) for JavaDoc and are followed by one or more parameters.

```
/*! \class Test
    \brief A test class.

    A more detailed description of class.
*/
```

In the above example the command `\class` is used. This indicates that the comment block contains documentation for the class 'Test'. Others are:

- o `\struct`: document a C-struct
- o `\union`: document a union
- o `\enum`: document an enumeration type
- o `\fn`: document a function
- o `\var`: document a variable, typedef, or enum value
- o `\def`: document a #define
- o `\typedef`: document a type definition
- o `\file`: document a file
- o `\namespace`: document a namespace
- o `\package`: document a Java package
- o `\interface`: document an IDL interface

3. Next, the contents of a special documentation block is parsed before being written to the HTML and / Latex output directories. This includes:
  1. Special commands are executed.
  2. Any white space and asterisks (\*) are removed.
  3. Blank lines are taken as new paragraphs.
  4. Words are linked to their corresponding documentation. Where the word is preceded by a percent sign (%) the percent sign is removed and the word remains.
  5. Where certain patterns are found in the text, links to members are created. Examples of this can be found on the [automatic link generation page](#) on the Doxygen documentation website.
  6. When the documentation is for Latex, HTML tags are interpreted and converted to Latex equivalents. A list of supported HTML tags can be found on the [HTML commands page](#) on the Doxygen documentation website.

### 6.1.5. Resources

More information can be found on the Doxygen website.

- [Doxygen homepage](#)
- [Doxygen introduction](#)
- [Doxygen documentation](#)
- [Output formats](#)

## APPENDIX A. APPENDIX

### A.1. MALLOPT

**mallocopt** is a library call that allows a program to change the behavior of the malloc memory allocator.

#### Example A.1. Allocator heuristics

An allocator has heuristics to determine long versus short lived objects. For the former, it attempts to allocate with **mmap**. For the later, it attempts to allocate with **sbrk**.

In order to override these heuristics, set **M\_MMAP\_THRESHOLD**.

In multi-threaded applications, the allocator creates multiple *arenas* in response to lock contention in existing arenas. This can improve the performance significantly for some multi-threaded applications at the cost of an increase in memory usage. To keep this under control, limit the number of arenas that can be created by using the **mallocopt** interface.

The allocator has limits on the number of arenas it can create. For 32bit targets, it will create  $2 * \#$  core arenas; for 64bit targets, it will create  $8 * \#$  core arenas. **mallocopt** allows the developer to override those limits.

#### Example A.2. mallocopt

To ensure no more than eight arenas are created, issue the following library call:

```
mallocopt (M_ARENA_MAX, 8);
```

The first argument for **mallocopt** can be:

- M\_MXFAST
- M\_TRIM\_THRESHOLD
- M\_TOP\_PAD
- M\_MMAP\_THRESHOLD
- M\_MMAP\_MAX
- M\_CHECK\_ACTION
- M\_PERTURB
- M\_ARENA\_TEST
- M\_ARENA\_MAX

Specific definitions for the above can be found at <http://www.makelinux.net/man/3/M/mallopt>.

**malloc\_trim**



**malloc\_trim** is a library call that requests the allocator return any unused memory back to the operating system. This is normally automatic when an object is freed. However, in some cases when freeing small objects, **glibc** might not immediately release the memory back to the operating system. It does this so that the free memory can be used to satisfy upcoming memory allocation requests as it is expensive to allocate from and release memory back to the operating system.

### **malloc\_stats**

**malloc\_stats** is used to dump information about the allocator's internal state to **stderr**. Using **mallinfo** is similar to this, but it places the state into a structure instead.

### **Further Information**

More information on **mallopt** can be found at <http://www.makelinux.net/man/3/M/mallopt> and [http://www.gnu.org/software/libc/manual/html\\_node/Malloc-Tunable-Parameters.html](http://www.gnu.org/software/libc/manual/html_node/Malloc-Tunable-Parameters.html).

## APPENDIX B. REVISION HISTORY

|  |                         |                        |
|--|-------------------------|------------------------|
| <b>Revision 6-9.3</b><br>Removal of section on distributed compiling.                        | <b>Thu 25 May 2017</b>  | <b>Vladimír Slávik</b> |
| <b>Revision 6-9.2</b><br>Release of the Developer Guide for Red Hat_Enterprise Linux 6.9     | <b>Mon 3 April 2017</b> | <b>Robert Krátký</b>   |
| <b>Revision 2-60</b><br>Release of the Developer Guide for Red Hat_Enterprise Linux 6.8      | <b>Wed 4 May 2016</b>   | <b>Robert Krátký</b>   |
| <b>Revision 2-56</b><br>Release of the Developer Guide for Red Hat_Enterprise Linux 6.7      | <b>Tue Jul 6 2015</b>   | <b>Robert Krátký</b>   |
| <b>Revision 2-55</b><br>Release of the Developer Guide for Red Hat_Enterprise Linux 6.7 Beta | <b>Wed Apr 15 2015</b>  | <b>Robert Krátký</b>   |
| <b>Revision 2-54</b><br>Update to sort order on the Red Hat Customer Portal.                 | <b>Tue Dec 16 2014</b>  | <b>Robert Krátký</b>   |
| <b>Revision 2-52</b><br>Re-release for RHSCCL 1.2 and DTS 3.0.                               | <b>Wed Nov 11 2014</b>  | <b>Robert Krátký</b>   |
| <b>Revision 2-51</b><br>Release of the Developer Guide for Red Hat_Enterprise Linux 6.6      | <b>Fri Oct 10 2014</b>  | <b>Robert Krátký</b>   |

---

# INDEX

## A

### advantages

- Python pretty-printers

  - debugging, [Python Pretty-Printers](#)

### Akonadi

- KDE Development Framework

  - libraries and runtime support, [KDE4 Architecture](#)

### architecture, KDE4

- KDE Development Framework

  - libraries and runtime support, [KDE4 Architecture](#)

### Autotools

- compiling and building, [Autotools](#)

## B

### backtrace

- tools

  - GNU debugger, [Simple GDB](#)

### Boost

- libraries and runtime support, [Boost](#)

### boost-doc

- Boost

  - libraries and runtime support, [Additional Information](#)

### breakpoint

- fundamentals

  - GNU debugger, [Simple GDB](#)

### breakpoints (conditional)

- GNU debugger, [Conditional Breakpoints](#)

### build-id

- compiling and building, [build-id Unique Identification of Binaries](#)

### building

compiling and building, [Compiling and Building](#)

## C

C++ Standard Library, GNU

libraries and runtime support, [The GNU C++ Standard Library](#)

cachegrind

tools

Valgrind, [Valgrind Tools](#)

callgrind

tools

Valgrind, [Valgrind Tools](#)

Collaborating, [Collaborating](#)

commands

fundamentals

GNU debugger, [Simple GDB](#)

profiling

Valgrind, [Valgrind Tools](#)

tools

Performance Counters for Linux (PCL) and perf, [Perf Tool Commands](#)

commonly-used commands

Autotools

compiling and building, [Autotools](#)

compatibility

libraries and runtime support, [Compatibility](#)

compiling a C Hello World program

usage

GCC, [Simple C Usage](#)

compiling a C++ Hello World program

usage

GCC, [Simple C++ Usage](#)

compiling and building

Autotools, [Autotools](#)

---

commonly-used commands, [Autotools](#)  
configuration script, [Configuration Script](#)  
documentation, [Autotools Documentation](#)  
plug-in for Eclipse, [Autotools Plug-in for Eclipse](#)  
templates (supported), [Autotools Plug-in for Eclipse](#)

build-id, [build-id Unique Identification of Binaries](#)  
GNU Compiler Collection, [GNU Compiler Collection \(GCC\)](#)  
documentation, [GCC Documentation](#)  
required packages, [Running GCC](#)  
usage, [Running GCC](#)

introduction, [Compiling and Building](#)

conditional breakpoints  
GNU debugger, [Conditional Breakpoints](#)

configuration script  
Autotools  
compiling and building, [Configuration Script](#)

continue  
tools  
GNU debugger, [Simple GDB](#)

## D

debugfs file system  
profiling  
ftrace, [ftrace](#)

debugging  
debuginfo-packages, [Installing Debuginfo Packages](#)  
installation, [Installing Debuginfo Packages](#)

GNU debugger, [GDB](#)  
fundamental mechanisms, [GDB](#)  
GDB, [GDB](#)  
requirements, [GDB](#)

introduction, [Debugging](#)  
Python pretty-printers, [Python Pretty-Printers](#)  
advantages, [Python Pretty-Printers](#)  
debugging output (formatted), [Python Pretty-Printers](#)

documentation, [Python Pretty-Printers](#)

pretty-printers, [Python Pretty-Printers](#)

variable tracking at assignments (VTA), [Variable Tracking at Assignments](#)

debugging a Hello World program

usage

GNU debugger, [Running GDB](#)

debugging output (formatted)

Python pretty-printers

debugging, [Python Pretty-Printers](#)

debuginfo-packages

debugging, [Installing Debuginfo Packages](#)

documentation

Autotools

compiling and building, [Autotools Documentation](#)

Boost

libraries and runtime support, [Additional Information](#)

GNU C++ Standard Library

libraries and runtime support, [Additional information](#)

GNU Compiler Collection

compiling and building, [GCC Documentation](#)

Java

libraries and runtime support, [Java Documentation](#)

KDE Development Framework

libraries and runtime support, [kdelibs Documentation](#)

OProfile

profiling, [OProfile Documentation](#)

Perl

libraries and runtime support, [Perl Documentation](#)

profiling

ftrace, [ftrace Documentation](#)

Python

libraries and runtime support, [Python Documentation](#)

---

**Python pretty-printers**

debugging, [Python Pretty-Printers](#)

**Qt**

libraries and runtime support, [Qt Library Documentation](#)

**Ruby**

libraries and runtime support, [Ruby Documentation](#)

**SystemTap**

profiling, [Additional Information](#)

**Valgrind**

profiling, [Additional information](#)

**Documentation****Doxygen, [Doxygen](#)**

Docment sources, [Documenting the Sources](#)

Getting Started, [Getting Started](#)

Resources, [Resources](#)

Running Doxygen, [Running Doxygen](#)

Supported output and languages, [Doxygen Supported Output and Languages](#)

**Documentation Tools, [Documentation Tools](#)****Doxygen****Documentation, [Doxygen](#)**

document sources, [Documenting the Sources](#)

Getting Started, [Getting Started](#)

Resources, [Resources](#)

Running Doxygen, [Running Doxygen](#)

Supported output and languages, [Doxygen Supported Output and Languages](#)

**E****execution (forked)**

GNU debugger, [Forked Execution](#)

**F****finish****tools**

GNU debugger, [Simple GDB](#)

**forked execution**

GNU debugger, [Forked Execution](#)

**formatted debugging output**

Python pretty-printers

debugging, [Python Pretty-Printers](#)

**framework (ftrace)**

profiling

ftrace, [ftrace](#)

**ftrace**

profiling, [ftrace](#)

debugfs file system, [ftrace](#)

documentation, [ftrace Documentation](#)

framework (ftrace), [ftrace](#)

usage, [Using ftrace](#)

**function tracer**

profiling

ftrace, [ftrace](#)

**fundamental commands**

fundamentals

GNU debugger, [Simple GDB](#)

**fundamental mechanisms**

GNU debugger

debugging, [GDB](#)

**fundamentals**

GNU debugger, [Simple GDB](#)

**G**

**gcc**

GNU Compiler Collection

compiling and building, [GNU Compiler Collection \(GCC\)](#)

**GCC C**

usage



---

compiling a C Hello World program, [Simple C Usage](#)

## **GCC C++**

usage

compiling a C++ Hello World program, [Simple C++ Usage](#)

## **GDB**

GNU debugger

debugging, [GDB](#)

## **Git**

configuration, [Installing and Configuring Git](#)

documentation, [Additional Resources](#)

installation, [Installing and Configuring Git](#)

overview, [Git](#)

usage, [Creating a New Repository](#)

## **GNOME Power Manager**

libraries and runtime support, [GNOME Power Manager](#)

## **gnome-power-manager**

GNOME Power Manager

libraries and runtime support, [GNOME Power Manager](#)

## **GNU C++ Standard Library**

libraries and runtime support, [The GNU C++ Standard Library](#)

## **GNU Compiler Collection**

compiling and building, [GNU Compiler Collection \(GCC\)](#)

## **GNU debugger**

conditional breakpoints, [Conditional Breakpoints](#)

debugging, [GDB](#)

execution (forked), [Forked Execution](#)

forked execution, [Forked Execution](#)

fundamentals, [Simple GDB](#)

breakpoint, [Simple GDB](#)

commands, [Simple GDB](#)

halting an executable, [Simple GDB](#)

inspecting the state of an executable, [Simple GDB](#)

starting an executable, [Simple GDB](#)

interfaces (CLI and machine), [Alternative User Interfaces for GDB](#)

thread and threaded debugging, [Debugging Individual Threads](#)

tools, [Simple GDB](#)

backtrace, [Simple GDB](#)

continue, [Simple GDB](#)

finish, [Simple GDB](#)

help, [Simple GDB](#)

list, [Simple GDB](#)

next, [Simple GDB](#)

print, [Simple GDB](#)

quit, [Simple GDB](#)

step, [Simple GDB](#)

usage, [Running GDB](#)

debugging a Hello World program, [Running GDB](#)

variations and environments, [Alternative User Interfaces for GDB](#)

## H

halting an executable

fundamentals

GNU debugger, [Simple GDB](#)

helgrind

tools

Valgrind, [Valgrind Tools](#)

help

tools

GNU debugger, [Simple GDB](#)

## I

inspecting the state of an executable

fundamentals

GNU debugger, [Simple GDB](#)

installation

debuginfo-packages

debugging, [Installing Debuginfo Packages](#)

interfaces (CLI and machine)

---

GNU debugger, [Alternative User Interfaces for GDB](#)

introduction

compiling and building, [Compiling and Building](#)

debugging, [Debugging](#)

libraries and runtime support, [Libraries and Runtime Support](#)

profiling, [Profiling](#)

SystemTap, [SystemTap](#)

ISO 14482 Standard C++ library

GNU C++ Standard Library

libraries and runtime support, [The GNU C++ Standard Library](#)

J

Java

libraries and runtime support, [Java](#)

K

KDE Development Framework

libraries and runtime support, [KDE Development Framework](#)

KDE4 architecture

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

kdelibs-devel

KDE Development Framework

libraries and runtime support, [KDE Development Framework](#)

kernel information packages

profiling

SystemTap, [SystemTap](#)

KHTML

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

KIO

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

**KJS****KDE Development Framework**libraries and runtime support, [KDE4 Architecture](#)**KNewStuff2****KDE Development Framework**libraries and runtime support, [KDE4 Architecture](#)**KXMLGUI****KDE Development Framework**libraries and runtime support, [KDE4 Architecture](#)**L****libraries**runtime support, [Libraries and Runtime Support](#)**libraries and runtime support****Boost**, [Boost](#)boost-doc, [Additional Information](#)documentation, [Additional Information](#)message passing interface (MPI), [Boost](#)meta-package, [Boost](#)**C++ Standard Library**, [GNU](#), [The GNU C++ Standard Library](#)compatibility, [Compatibility](#)**GNOME Power Manager**, [GNOME Power Manager](#)gnome-power-manager, [GNOME Power Manager](#)**GNU C++ Standard Library**, [The GNU C++ Standard Library](#)documentation, [Additional information](#)ISO 14482 Standard C++ library, [The GNU C++ Standard Library](#)libstdc++-devel, [The GNU C++ Standard Library](#)libstdc++-docs, [Additional information](#)Standard Template Library, [The GNU C++ Standard Library](#)introduction, [Libraries and Runtime Support](#)**Java**, [Java](#)documentation, [Java Documentation](#)**KDE Development Framework**, [KDE Development Framework](#)Akonadi, [KDE4 Architecture](#)documentation, [kdelibs Documentation](#)

KDE4 architecture, [KDE4 Architecture](#)  
kdelibs-devel, [KDE Development Framework](#)  
KHTML, [KDE4 Architecture](#)  
KIO, [KDE4 Architecture](#)  
KJS, [KDE4 Architecture](#)  
KNewStuff2, [KDE4 Architecture](#)  
KXMLGUI, [KDE4 Architecture](#)  
Phonon, [KDE4 Architecture](#)  
Plasma, [KDE4 Architecture](#)  
Solid, [KDE4 Architecture](#)  
Sonnet, [KDE4 Architecture](#)  
Strigi, [KDE4 Architecture](#)  
Telepathy, [KDE4 Architecture](#)

libstdc++, [The GNU C++ Standard Library](#)

Perl, [Perl](#)

documentation, [Perl Documentation](#)  
module installation, [Installation](#)  
updates, [Perl Updates](#)

Python, [Python](#)

documentation, [Python Documentation](#)  
updates, [Python Updates](#)

Qt, [Qt](#)

documentation, [Qt Library Documentation](#)  
meta object compiler (MOC), [Qt](#)  
Qt Creator, [Qt Creator](#)  
qt-doc, [Qt Library Documentation](#)  
updates, [Qt Updates](#)  
widget toolkit, [Qt](#)

Ruby, [Ruby](#)

documentation, [Ruby Documentation](#)  
ruby-devel, [Ruby](#)

## Library and Runtime Details

NSS Shared Databases, [NSS Shared Databases](#)

Backwards Compatibility, [Backwards Compatibility](#)

Documentation, [NSS Shared Databases Documentation](#)

## libstdc++

libraries and runtime support, [The GNU C++ Standard Library](#)

**libstdc++-devel**

GNU C++ Standard Library

libraries and runtime support, [The GNU C++ Standard Library](#)**libstdc++-docs**

GNU C++ Standard Library

libraries and runtime support, [Additional information](#)**list**

tools

GNU debugger, [Simple GDB](#)Performance Counters for Linux (PCL) and perf, [Perf Tool Commands](#)**M****machine interface**GNU debugger, [Alternative User Interfaces for GDB](#)**malloc, [malloc](#)****massif**

tools

Valgrind, [Valgrind Tools](#)**mechanisms**

GNU debugger

debugging, [GDB](#)**memcheck**

tools

Valgrind, [Valgrind Tools](#)**message passing interface (MPI)**

Boost

libraries and runtime support, [Boost](#)**meta object compiler (MOC)**

Qt

libraries and runtime support, [Qt](#)**meta-package**

Boost

---

libraries and runtime support, [Boost](#)

module installation

Perl

libraries and runtime support, [Installation](#)

**N**

next

tools

GNU debugger, [Simple GDB](#)

**NSS Shared Databases**

Library and Runtime Details, [NSS Shared Databases](#)

Backwards Compatibility, [Backwards Compatibility](#)

Documentation, [NSS Shared Databases Documentation](#)

**O**

**OProfile**

profiling, [OProfile](#)

documentation, [OProfile Documentation](#)

usage, [Using OProfile](#)

**P**

**perf**

profiling

Performance Counters for Linux (PCL) and perf, [Performance Counters for Linux \(PCL\) Tools and perf](#)

usage

Performance Counters for Linux (PCL) and perf, [Using Perf](#)

**Performance Counters for Linux (PCL) and perf**

profiling, [Performance Counters for Linux \(PCL\) Tools and perf](#)

subsystem (PCL), [Performance Counters for Linux \(PCL\) Tools and perf](#)

tools, [Perf Tool Commands](#)

commands, [Perf Tool Commands](#)

list, [Perf Tool Commands](#)

record, [Perf Tool Commands](#)

report, [Perf Tool Commands](#)

stat, [Perf Tool Commands](#)

usage, [Using Perf](#)

perf, [Using Perf](#)

## Perl

libraries and runtime support, [Perl](#)

## Phonon

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

## Plasma

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

## plug-in for Eclipse

Autotools

compiling and building, [Autotools Plug-in for Eclipse](#)

## pretty-printers

Python pretty-printers

debugging, [Python Pretty-Printers](#)

## print

tools

GNU debugger, [Simple GDB](#)

## profiling

conflict between perf and oprofile, [Using Perf](#)

ftrace, [ftrace](#)

introduction, [Profiling](#)

OProfile, [OProfile](#)

Performance Counters for Linux (PCL) and perf, [Performance Counters for Linux \(PCL\) Tools and perf](#)

SystemTap, [SystemTap](#)

Valgrind, [Valgrind](#)

## Python

libraries and runtime support, [Python](#)

## Python pretty-printers



debugging, [Python Pretty-Printers](#)

## Q

### Qt

libraries and runtime support, [Qt](#)

### Qt Creator

#### Qt

libraries and runtime support, [Qt Creator](#)

### qt-doc

#### Qt

libraries and runtime support, [Qt Library Documentation](#)

### quit

#### tools

GNU debugger, [Simple GDB](#)

## R

### record

#### tools

Performance Counters for Linux (PCL) and perf, [Perf Tool Commands](#)

### report

#### tools

Performance Counters for Linux (PCL) and perf, [Perf Tool Commands](#)

### required packages

#### GNU Compiler Collection

compiling and building, [Running GCC](#)

#### profiling

SystemTap, [SystemTap](#)

### requirements

#### GNU debugger

debugging, [GDB](#)

### Revision control, [Collaborating](#)

### Ruby

libraries and runtime support, [Ruby](#)

ruby-devel

Ruby

libraries and runtime support, [Ruby](#)

runtime support

libraries, [Libraries and Runtime Support](#)

## S

scripts (SystemTap scripts)

profiling

SystemTap, [SystemTap](#)

Solid

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

Sonnet

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

Standard Template Library

GNU C++ Standard Library

libraries and runtime support, [The GNU C++ Standard Library](#)

starting an executable

fundamentals

GNU debugger, [Simple GDB](#)

stat

tools

Performance Counters for Linux (PCL) and perf, [Perf Tool Commands](#)

step

tools

GNU debugger, [Simple GDB](#)

Strigi

KDE Development Framework

---

libraries and runtime support, [KDE4 Architecture](#)

subsystem (PCL)

profiling

Performance Counters for Linux (PCL) and perf, [Performance Counters for Linux \(PCL\) Tools and perf](#)

supported templates

Autotools

compiling and building, [Autotools Plug-in for Eclipse](#)

SystemTap

profiling, [SystemTap](#)

documentation, [Additional Information](#)

introduction, [SystemTap](#)

kernel information packages, [SystemTap](#)

required packages, [SystemTap](#)

scripts (SystemTap scripts), [SystemTap](#)

T

Telepathy

KDE Development Framework

libraries and runtime support, [KDE4 Architecture](#)

templates (supported)

Autotools

compiling and building, [Autotools Plug-in for Eclipse](#)

thread and threaded debugging

GNU debugger, [Debugging Individual Threads](#)

tools

GNU debugger, [Simple GDB](#)

Performance Counters for Linux (PCL) and perf, [Perf Tool Commands](#)

profiling

Valgrind, [Valgrind Tools](#)

Valgrind, [Valgrind Tools](#)

U

**updates****Perl**

libraries and runtime support, [Perl Updates](#)

**Python**

libraries and runtime support, [Python Updates](#)

**Qt**

libraries and runtime support, [Qt Updates](#)

**usage****GNU Compiler Collection**

compiling and building, [Running GCC](#)

**GNU debugger, [Running GDB](#)**

fundamentals, [Simple GDB](#)

**Performance Counters for Linux (PCL) and perf, [Using Perf](#)**  
profiling

ftrace, [Using ftrace](#)

OProfile, [Using OProfile](#)

**Valgrind**

profiling, [Using Valgrind](#)

**V****Valgrind**

profiling, [Valgrind](#)

commands, [Valgrind Tools](#)

documentation, [Additional information](#)

tools, [Valgrind Tools](#)

usage, [Using Valgrind](#)

**tools**

cachegrind, [Valgrind Tools](#)

callgrind, [Valgrind Tools](#)

helgrind, [Valgrind Tools](#)

massif, [Valgrind Tools](#)

memcheck, [Valgrind Tools](#)

**variable tracking at assignments (VTA)**

debugging, [Variable Tracking at Assignments](#)

variations and environments

GNU debugger, [Alternative User Interfaces for GDB](#)

Version control, [Collaborating](#)

W

widget toolkit

Qt

libraries and runtime support, [Qt](#)