

# Red Hat Developer Toolset 1.1

## User Guide

### Installing and Using Red Hat Developer Toolset

#### Edition 1



**Jaromír Hradílek**

Red Hat Engineering Content Services

[jhradilek@redhat.com](mailto:jhradilek@redhat.com)

**Jacquelynn East**

Red Hat Engineering Content Services

[jeast@redhat.com](mailto:jeast@redhat.com)

**Matt Newsome**

Red Hat Software Engineering

[mnewsome@redhat.com](mailto:mnewsome@redhat.com)

## Legal Notice

Copyright © 2012, 2013 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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, MetaMatrix, 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.

All other trademarks are the property of their respective owners.

1801 Varsity Drive  
Raleigh, NC 27606-2072 USA  
Phone: +1 919 754 3700  
Phone: 888 733 4281  
Fax: +1 919 754 3701

## Abstract

Red Hat Developer Toolset is a Red Hat offering for developers on the Red Hat Enterprise Linux platform. The *Red Hat Developer Toolset User Guide* provides an overview of this product, explains how to invoke and use the Developer Toolset versions of the tools, and links to resources with more in-depth information.

# Table of Contents

## Preface

1. Document Conventions
  - 1.1. Typographic Conventions
  - 1.2. Pull-quote Conventions
  - 1.3. Notes and Warnings
2. Getting Help and Giving Feedback
  - 2.1. Do You Need Help?
  - 2.2. We Need Feedback!

## I. Introduction

1. Red Hat Developer Toolset
  - 1.1. About Red Hat Developer Toolset
  - 1.2. Main Features
  - 1.3. Compatibility
  - 1.4. Getting Access to Red Hat Developer Toolset
    - 1.4.1. Using RHN Classic
    - 1.4.2. Using Red Hat Subscription Management
  - 1.5. Installing Red Hat Developer Toolset
  - 1.6. Updating Red Hat Developer Toolset
  - 1.7. Uninstalling Red Hat Developer Toolset
  - 1.8. Additional Resources

## II. Development Tools

2. GNU Compiler Collection (GCC)
  - 2.1. GNU C Compiler
    - 2.1.1. Installing the C Compiler
    - 2.1.2. Using the C Compiler
    - 2.1.3. Running a C Program
  - 2.2. GNU C++ Compiler
    - 2.2.1. Installing the C++ Compiler
    - 2.2.2. Using the C++ Compiler
    - 2.2.3. Running a C++ Program
  - 2.3. GNU Fortran Compiler
    - 2.3.1. Installing the Fortran Compiler
    - 2.3.2. Using the Fortran Compiler
    - 2.3.3. Running a Fortran Program
  - 2.4. Additional Resources
3. GNU Debugger (GDB)
  - 3.1. Installing the GNU Debugger
  - 3.2. Preparing a Program for Debugging
  - 3.3. Running the GNU Debugger
  - 3.4. Listing Source Code
  - 3.5. Setting Breakpoints
  - 3.6. Starting Execution
  - 3.7. Displaying Current Values
  - 3.8. Continuing Execution
  - 3.9. Additional Resources

#### 4. binutils

- 4.1. Installing binutils
- 4.2. Using the GNU Assembler
- 4.3. Using the GNU Linker
- 4.4. Using Other Binary Tools
- 4.5. Additional Resources

#### 5. elfutils

- 5.1. Installing elfutils
- 5.2. Using elfutils
- 5.3. Additional Resources

#### 6. dwz

- 6.1. Installing dwz
- 6.2. Using dwz
- 6.3. Additional Resources

### III. Performance Monitoring Tools

#### 7. SystemTap

- 7.1. Installing SystemTap
- 7.2. Using SystemTap
- 7.3. Additional Resources

#### 8. Valgrind

- 8.1. Installing Valgrind
- 8.2. Using Valgrind
- 8.3. Additional Resources

#### 9. OProfile

- 9.1. Installing OProfile
- 9.2. Using OProfile
- 9.3. Additional Resources

### IV. Integrated Development Environments

#### 10. Eclipse

- 10.1. Installing Eclipse
- 10.2. Running Eclipse
- 10.3. Additional Resources

### V. Getting Help

#### 11. Accessing Red Hat Product Documentation

#### 12. Accessing the Customer Portal

- 12.1. The Plan Menu
- 12.2. The Deploy Menu
- 12.3. The Connect Menu

#### 13. Contacting Global Support Services

- 13.1. Gathering Required Information
- 13.2. Escalating an Issue
- 13.3. Re-opening a Service Request
- 13.4. Additional Resources

### A. Changes in Version 1.0

#### A.1. Changes in GCC 4.7.0

- A.1.1. Status and Features
- A.1.2. Language Compatibility
- A.1.3. ABI Compatibility
- A.1.4. Debugging Compatibility
- A.1.5. Other Compatibility

- A.2. Changes in GDB 7.4
  - A.2.1. New Features
  - A.2.2. Compatibility Changes

- A.3. Changes in binutils 2.22.52
  - A.3.1. GNU Linker
  - A.3.2. GNU Assembler
  - A.3.3. Other Binary Tools

## B. Changes in Version 1.1

- B.1. Changes in GCC 4.7.2
  - B.1.1. C++11 Compatibility
  - B.1.2. Fortran
  - B.1.3. Architecture-specific Options

- B.2. Changes in GDB 7.5
- B.3. Changes in binutils 2.23.51
  - B.3.1. GNU Linker
  - B.3.2. GNU Assembler
  - B.3.3. Other Binary Tools

- B.4. Changes in elfutils 0.154
- B.5. Changes in Valgrind 3.8.1
- B.6. Changes in OProfile 0.154

## C. Revision History

Index

# Preface

## 1. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the [Liberation Fonts](#) set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

### 1.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

#### Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keys and key combinations. For example:

To see the contents of the file **my\_next\_bestselling\_novel** in your current working directory, enter the **cat my\_next\_bestselling\_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a key, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from an individual key by the plus sign that connects each part of a key combination. For example:

Press **Enter** to execute the command.

Press **Ctrl+Alt+F2** to switch to a virtual terminal.

The first example highlights a particular key to press. The second example highlights a key combination: a set of three keys pressed simultaneously.

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **mono-spaced bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

#### Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog box text; labeled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications** → **Accessories** → **Character Map** from the main menu bar. Next, choose **Search** → **Find...** from the

**Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

### *Mono-spaced Bold Italic* or *Proportional Bold Italic*

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh john@example.com**.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount /home**.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — username, domain.name, file-system, package, version and release. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, italics denotes the first use of a new and important term. For example:

Publican is a *DocBook* publishing system.

## 1.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads      images  notes  scripts svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:



```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref     = iniCtx.lookup("EchoBean");
        EchoHome        home    = (EchoHome) ref;
        Echo             echo    = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

### 1.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



#### Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



#### Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' will not cause data loss but may cause irritation and frustration.



#### Warning

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

## 2. Getting Help and Giving Feedback

### 2.1. Do You Need Help?

If you experience difficulty with a procedure described in this documentation, visit the Red Hat Customer Portal at <http://access.redhat.com>. Through the customer portal, you can:

- ▶ search or browse through a knowledgebase of technical support articles about Red Hat products.
- ▶ submit a support case to Red Hat Global Support Services (GSS).
- ▶ access other product documentation.

Red Hat also hosts a large number of electronic mailing lists for discussion of Red Hat software and

technology. You can find a list of publicly available mailing lists at <https://www.redhat.com/mailman/listinfo>. Click on the name of any mailing list to subscribe to that list or to access the list archives.

## 2.2. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: <http://bugzilla.redhat.com/> against the product **Red Hat Developer Toolset**.

When submitting a bug report, be sure to mention the manual's identifier: *doc-User\_Guide*

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, please include the section number and some of the surrounding text so we can find it easily.

---

## Part I. Introduction

# Chapter 1. Red Hat Developer Toolset

## 1.1. About Red Hat Developer Toolset

**Red Hat Developer Toolset** is a Red Hat offering for developers on the Red Hat Enterprise Linux platform, and provides a complete set of development and performance analysis tools that can be installed and used on multiple versions of Red Hat Enterprise Linux. Executables built with the Red Hat Developer Toolset toolchain can then also be deployed and run on multiple versions of Red Hat Enterprise Linux. For detailed compatibility information, refer to [Section 1.3, “Compatibility”](#).

Critically, Red Hat Developer Toolset does not replace the default system tools provided with Red Hat Enterprise Linux 5 or 6 when installed on those platforms. Instead, a parallel set of developer tools provides an alternative, newer version of those tools for optional use by developers. The default compiler and debugger, for example, remain those provided by the base Red Hat Enterprise Linux system.

Red Hat Developer Toolset differs from “[Technology Preview](#)” compiler releases previously supplied in Red Hat Enterprise Linux in two important respects:

- ▶ Red Hat Developer Toolset can be used on multiple major and minor releases of Red Hat Enterprise Linux, as detailed in [Section 1.3, “Compatibility”](#).
- ▶ Unlike Technology Preview compilers and other tools shipped in earlier Red Hat Enterprise Linux, Red Hat Developer Toolset is fully supported under Red Hat Enterprise Linux Subscription Level Agreements, is functionally complete, and is intended for production use.

Important bug fixes and security errata will be issued to Red Hat Developer Toolset subscribers in a similar manner to Red Hat Enterprise Linux for two years from the release of each major version release. New major versions of Red Hat Developer Toolset will be released annually, providing significant updates for existing components and adding major new components. A single minor release, issued six months after each new major version release, will provide a smaller update of bug fixes, security errata, and new minor components.

Additionally, the Red Hat Enterprise Linux Application Compatibility Specification also applies to Red Hat Developer Toolset (subject to some constraints on the use of newer C++11 language features, detailed in [Section A.1.3, “ABI Compatibility”](#)).



### Important

Applications and libraries provided by Red Hat Developer Toolset do not replace the Red Hat Enterprise Linux system versions, nor are they used in preference to the system versions. Using a framework called **Software Collections**, an additional set of developer tools is installed into the `/opt` directory and is explicitly enabled by the user on demand using the supplied `sc1` utility.

## 1.2. Main Features

Red Hat Developer Toolset 1.1 provides current versions of the following tools:

- ▶ **GNU Compiler Collection (GCC)** version **4.7.2**;
- ▶ **GNU Debugger (GDB)** version **7.5**;
- ▶ **binutils** version **2.23.51**;
- ▶ **elfutils** version **0.154**;
- ▶ **dwz** version **0.7**;

- ▶ **SystemTap** version **1.8**;
- ▶ **Valgrind** version **3.8.1**;
- ▶ **OProfile** version **0.9.7**.

In particular, the Red Hat Developer Toolset version of the **GNU Compiler Collection (GCC)** provides the following features:

- ▶ Support for the Fortran programming language.
- ▶ Experimental support for the C++11 standard.
- ▶ Experimental support for C++11 atomic types and Transactional Memory.
- ▶ Improved support for link-time optimization (LTO).
- ▶ Improved support for interprocedural optimization.
- ▶ New **-Ofast** general optimization level.
- ▶ New string length optimization pass.
- ▶ Various compile time and memory usage improvements.
- ▶ Support for OpenMP 3.1, an API specification for parallel programming.
- ▶ Optimization for various new Intel and AMD processors.

The version of the **GNU Debugger (GDB)** included in Red Hat Developer Toolset provides the following features:

- ▶ Improved and expanded support for Python scripting.
- ▶ Improved handling of C++ debuggee executables.
- ▶ Improved inferior control commands.
- ▶ Improved support for ambiguous line specifications.
- ▶ Improved tracepoint support.
- ▶ Multi-program debugging.

Additionally, the Red Hat Developer Toolset version of **binutils** provides these features:

- ▶ The new **gold** linker, which is smaller and faster than **ld**. Note that **gold** is not the default linker and must be explicitly enabled by using the **alternatives** command.
- ▶ Support for link-time optimization (LTO) in conjunction with GCC.
- ▶ Support for build-IDs, unique numbers to identify executables.
- ▶ Support for the **IFUNC** and **UNIQUE** symbols that are used by **glibc** to improve performance. Due to dependencies on a particular version of the **glibc** library, these symbols are only available on Red Hat Enterprise Linux 6.
- ▶ Compressed debug sections for smaller debug info files.

For a full list of changes and features introduced in this release, refer to [Appendix B, Changes in Version 1.1](#).

## 1.3. Compatibility

Red Hat Developer Toolset 1.1 is available for Red Hat Enterprise Linux 5 and 6, both for 32-bit and 64-bit Intel and AMD architectures. [Table 1.1, “Red Hat Developer Toolset 1.1 Compatibility Matrix”](#) illustrates the support for binaries built with Red Hat Developer Toolset on a certain version of Red Hat Enterprise Linux when those binaries are run on various other versions of this system.

For ABI compatibility information, refer to [Section A.1.3, “ABI Compatibility”](#).

**Table 1.1. Red Hat Developer Toolset 1.1 Compatibility Matrix**

		“Run on”									
		< 5.6	5.6 EUS	5.7	5.8	5.9	6.0 EUS	6.1 EUS	6.2 EUS	6.3	6.4
“Build with”	< 5.6	Unsupported version of the Red Hat Enterprise Linux host									
	5.6 EUS	✗	✓	✗	✗	✓	✗	✓	✓	✓	✓
	5.7	Unsupported version of the Red Hat Enterprise Linux host									
	5.8	Unsupported version of the Red Hat Enterprise Linux host									
	5.9	✗	✗	✗	✗	✓	✗	✓	✓	✓	✓
	6.0 EUS	Unsupported version of the Red Hat Enterprise Linux host									
	6.1 EUS	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓
	6.2 EUS	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓
	6.3	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓
	6.4	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓

✗ Unsupported

✓ Supported

## 1.4. Getting Access to Red Hat Developer Toolset

Red Hat Developer Toolset is an offering that is distributed as part of the developer subscriptions listed in [Table 1.2, “Subscriptions that provide Red Hat Developer Toolset”](#). Depending on the subscription management service with which you registered your Red Hat Enterprise Linux system, you can either enable Red Hat Developer Toolset by using the Red Hat Subscription Management, or by using RHN Classic.

For detailed instructions on how to enable Red Hat Developer Toolset using RHN Classic or Red Hat Subscription Management, refer to the respective section below. For information on how to register your system with one of these subscription management services, see the *Red Hat Subscription Management Guide*.



### Important

If you are running a version of Red Hat Enterprise Linux prior to 5.9 or 6.4, you will be unable to download Red Hat Developer Toolset through Red Hat Subscription Management. To obtain Red Hat Developer Toolset, you will need to either update to Red Hat Enterprise Linux 5.9 or 6.4, or register your system with RHN Classic. For more information, refer to <https://access.redhat.com/knowledge/solutions/129003>.

**Table 1.2. Subscriptions that provide Red Hat Developer Toolset**

SKU	Subscription Name
RH2282403	Red Hat Enterprise Linux Developer Support, Professional
RH2264664	Red Hat Enterprise Linux Developer Support, Enterprise
RH2262474	Red Hat Enterprise Linux Developer Suite
RH3482727	Red Hat Enterprise Linux Developer Workstation, Professional
RH3413336	Red Hat Enterprise Linux Developer Workstation, Enterprise
RH3414706	30 day Self-Supported Red Hat Enterprise Linux Developer Workstation Evaluation
RH3474212	60 day Supported Red Hat Enterprise Linux Developer Workstation Evaluation
RH3437268	90 day Supported Red Hat Enterprise Linux Developer Workstation Evaluation
SER0402	1-year Unsupported Partner Evaluation Red Hat Enterprise Linux
SER0403	1-year Unsupported Red Hat Advanced Partner Subscription

### 1.4.1. Using RHN Classic

If your system is registered with RHN Classic, complete the following steps to subscribe to Red Hat Developer Toolset:

1. Display a list of all channels that are available to you in order to determine the exact name of the Red Hat Developer Toolset channel. To do so, type the following at a shell prompt as **root**:

```
rhn-channel --available-channels
```

The name of the channel depends on the specific version of Red Hat Enterprise Linux you are using and is in the **rhel-architecture-variant-dts-version** format, where **architecture** is the system's CPU architecture (**x86\_64** or **i386**), **variant** is the Red Hat Enterprise Linux system variant (**server** or **workstation**), and **version** is the Red Hat Enterprise Linux system version (**5** or **6**).

2. Subscribe the system to the Red Hat Developer Toolset channel by running the following command as **root**:

```
rhn-channel --add --channel=channel_name
```

Replace **channel\_name** with the name you determined in the previous step.

3. To verify the list of channels you are subscribed to, at any time, run as **root**:

```
rhn-channel --list
```

Once the system is subscribed, you can install Red Hat Developer Toolset as described in [Section 1.5, "Installing Red Hat Developer Toolset"](#). For more information on how to register your system with RHN Classic, refer to the *Red Hat Subscription Management Guide*.

### 1.4.2. Using Red Hat Subscription Management

If your system is registered with Red Hat Subscription Management, complete the following steps to attach a subscription that provides access to the repository for Red Hat Developer Toolset, and then enable that repository:

1. Display a list of all subscriptions that are available for your system to determine the pool ID of a subscription that provides Red Hat Developer Toolset. To do so, type the following at a shell prompt as **root**:

```
subscription-manager list --available
```

For each available subscription, this command displays its name, unique identifier, expiration date, and other details related to your subscription. The pool ID is listed on a line beginning with **Pool Id**.

For a complete list of subscriptions that provide access to Red Hat Developer Toolset, refer to [Table 1.2, “Subscriptions that provide Red Hat Developer Toolset”](#).

2. Attach the appropriate subscription to your system by running the following command as **root**:

```
subscription-manager subscribe --pool=pool_id
```

Replace *pool\_id* with the pool ID you determined in the previous step. To verify the list of subscriptions your system has currently attached, at any time, run as **root**:

```
subscription-manager list --consumed
```

3. Display a list of available Yum repositories to retrieve repository metadata and to determine the exact name of the Red Hat Developer Toolset repositories. As **root**, type:

```
yum repolist all
```

The repository names depend on the specific version of Red Hat Enterprise Linux you are using, and are in the following format:

```
rhel-variant-dts-version-rpms
rhel-variant-dts-version-debug-rpms
rhel-variant-dts-version-source-rpms
```

Replace *variant* with the Red Hat Enterprise Linux system variant (**server** or **workstation**), and *version* with the Red Hat Enterprise Linux system version (**5** or **6**).

4. Enable the appropriate repository. On Red Hat Enterprise Linux 6, you can do so by running the following command as **root**:

```
yum-config-manager --enable repository
```

On Red Hat Enterprise Linux 5, which does not support the **yum-config-manager** tool, edit the `/etc/yum.repos.d/redhat.repo` file, locate the relevant [*repository*] section, and set the value of the **enabled** option to **1**:

```
enabled = 1
```

Once the subscription is attached to the system, you can install Red Hat Developer Toolset as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#). For more information on how to register your system using Red Hat Subscription Management and associate it with subscriptions, refer to the *Red Hat Subscription Management Guide*.

## 1.5. Installing Red Hat Developer Toolset

Red Hat Developer Toolset is distributed as a collection of RPM packages. To install the tools and libraries provided by Red Hat Developer Toolset, ensure that the system is associated with an appropriate subscription (see [Section 1.4, “Getting Access to Red Hat Developer Toolset”](#)), and install the `devtoolset-1.1` package by typing the following at a shell prompt as **root**:



```
yum install devtoolset-1.1
```

This installs the GNU Compiler Collection, GNU Debugger, binutils, the **scl** utility, and other dependent packages to the system. You can also install any of the additional packages listed in the *Package Manifest* section of the *Red Hat Developer Toolset 1.1 Release Notes* by running the command in the following form as **root**:

```
yum install package_name
```

For example, to install *libitm* and *devtoolset-1.1-gdb-gdbserver*, type:

```
yum install libitm devtoolset-1.1-gdb-gdbserver
```

To install debugging information for any of the Red Hat Developer Toolset packages, run as **root**:

```
debuginfo-install package_name
```

## 1.6. Updating Red Hat Developer Toolset

Before updating an existing installation of Red Hat Developer Toolset on your system, make sure the *devtoolset-1.0-gcc-debuginfo* package is not installed. To do so, type the following at a shell prompt as **root**:

```
yum remove devtoolset-1.0-gcc-debuginfo
```

Then run the following command as **root** to update your Red Hat Enterprise Linux installation:

```
yum update
```

This updates all packages on your Red Hat Enterprise Linux system, including the Red Hat Developer Toolset versions of the GNU Compiler Collection, GNU Debugger, binutils, and other dependent packages. Note that if you need additional packages with debugging information for any of the Red Hat Developer Toolset packages (such as *devtoolset-1.1-gcc-debuginfo*), you can install them at any time by using the following command as **root**:

```
debuginfo-install package_name
```

## 1.7. Uninstalling Red Hat Developer Toolset

To uninstall Red Hat Developer Toolset packages from your system, type the following at a shell prompt as **root**:

```
yum remove devtoolset-1.1\*
```

This removes the GNU Compiler Collection, GNU Debugger, binutils, and other packages that are part of Red Hat Developer Toolset from the system. To uninstall the **scl** utility as well, type as **root**:

```
yum remove scl-utils\*
```

Note that uninstallation of the tools provided by Red Hat Developer Toolset does not affect the Red Hat Enterprise Linux system versions of these tools.

## 1.8. Additional Resources

For more information about Red Hat Developer Toolset and Red Hat Enterprise Linux, refer to the resources listed below.

### Online Documentation

- ▶ [Red Hat Subscription Management Guide](#) — The *Red Hat Subscription Management Guide* provides detailed information on how to manage subscriptions on Red Hat Enterprise Linux.
- ▶ [Red Hat Developer Toolset 1.1 Release Notes](#) — The *Release Notes* for Red Hat Developer Toolset 1.1 contain more information about this product.
- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on the Eclipse IDE, libraries and runtime support, compiling and building, debugging, and profiling on this system.
- ▶ [Red Hat Enterprise Linux 6 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 6 explains how to obtain, install, and update the system.
- ▶ [Red Hat Enterprise Linux 5 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 5 explains how to obtain, install, and update the system.
- ▶ [Red Hat Enterprise Linux 6 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 6 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 6.
- ▶ [Red Hat Enterprise Linux 5 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 5 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 5.

### See Also

- ▶ [Appendix B, Changes in Version 1.1](#) provides a comprehensive list of changes and improvements over the Red Hat Enterprise Linux system versions of the GNU Compiler Collection, GNU Debugger, and binutils, as well as information about the language, ABI, and debugging compatibility.

---

## Part II. Development Tools

## Chapter 2. GNU Compiler Collection (GCC)

The **GNU Compiler Collection**, commonly abbreviated **GCC**, is a portable compiler suite with support for a wide selection of programming languages.

Red Hat Developer Toolset is distributed with **GCC 4.7.2**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements, including optimization for various new Intel and AMD processors, support for OpenMP 3.1 and link-time optimization. This version also includes experimental support for the C++11 standard, C++11 atomic types, and Transactional Memory. For a detailed list of changes, refer to [Section B.1, “Changes in GCC 4.7.2”](#).

### 2.1. GNU C Compiler

#### 2.1.1. Installing the C Compiler

In Red Hat Developer Toolset, the GNU C compiler is provided by the *devtoolset-1.1-gcc* package, and is automatically installed with *devtoolset-1.1* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

#### 2.1.2. Using the C Compiler

To compile a C program on the command line, run the **gcc** compiler as follows:

```
scl enable devtoolset-1.1 'gcc -o output_file source_file...'
```

This creates a binary file named *output\_file* in the current working directory. If the **-o** option is omitted, the compiler creates a file named **a.out** by default.

When you are working on a project that consists of several source files, it is common to compile an object file for each of the source files first and then link these object files together. This way, when you change a single source file, you can recompile only this file without having to compile the entire project. To compile an object file on the command line, run the following command:

```
scl enable devtoolset-1.1 'gcc -o object_file -c source_file'
```

This creates an object file named *object\_file*. If the **-o** option is omitted, the compiler creates a file named after the source file with the **.o** file extension. To link object files together and create a binary file, run:

```
scl enable devtoolset-1.1 'gcc -o output_file object_file...'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **gcc** as default:

```
scl enable devtoolset-1.1 'bash'
```



## Note

To verify the version of **gcc** you are using at any point, type the following at a shell prompt:

```
which gcc
```

Red Hat Developer Toolset's **gcc** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **gcc**:

```
gcc -v
```



## Important

Some newer library features are statically linked into applications built with Red Hat Developer Toolset to support execution on multiple versions of Red Hat Enterprise Linux. This adds a small additional security risk as normal Red Hat Enterprise Linux errata would not change this code. If the need for developers to rebuild their applications due to such an issue arises, Red Hat will signal this via a security erratum. Developers are strongly advised not to statically link their entire application for the same reasons.

### Example 2.1. Compiling a C Program on the Command Line

Imagine that you have a source file named **hello.c** with the following contents:

```
#include <stdio.h>

int main(int argc, char *argv[]) {
    printf("Hello, World!\n");
    return 0;
}
```

To compile this source code on the command line by using the **gcc** compiler from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-1.1 'gcc -o hello hello.c'
```

This creates a new binary file called **hello** in the current working directory.

### 2.1.3. Running a C Program

When **gcc** compiles a program, it creates an executable binary file. To run this program on the command line, change to the directory with the executable file and type:

```
./file_name
```

### Example 2.2. Running a C Program on the Command Line

Assuming that you have successfully compiled the `hello` binary file as shown in [Example 2.1, “Compiling a C Program on the Command Line”](#), you can run it by typing the following at a shell prompt:

```
~]$ ./hello
Hello, World!
```

## 2.2. GNU C++ Compiler

### 2.2.1. Installing the C++ Compiler

In Red Hat Developer Toolset, the GNU C++ compiler is provided by the `devtoolset-1.1-gcc-c++` package, and is automatically installed with the `devtoolset-1.1` package as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 2.2.2. Using the C++ Compiler

To compile a C++ program on the command line, run the `g++` compiler as follows:

```
sc1 enable devtoolset-1.1 'g++ -o output_file source_file...'
```

This creates a binary file named `output_file` in the current working directory. If the `-o` option is omitted, the `g++` compiler creates a file named `a.out` by default.

When you are working on a project that consists of several source files, it is common to compile an object file for each of the source files first and then link these object files together. This way, when you change a single source file, you can recompile only this file without having to compile the entire project. To compile an object file on the command line, run the following command:

```
sc1 enable devtoolset-1.1 'g++ -o object_file -c source_file'
```

This creates an object file named `object_file`. If the `-o` option is omitted, the `g++` compiler creates a file named after the source file with the `.o` file extension. To link object files together and create a binary file, run:

```
sc1 enable devtoolset-1.1 'g++ -o output_file object_file...'
```

Note that you can execute any command using the `sc1` utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset `g++` as default:

```
sc1 enable devtoolset-1.1 'bash'
```

 Note

To verify the version of **g++** you are using at any point, type the following at a shell prompt:

```
which g++
```

Red Hat Developer Toolset's **g++** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **g++**:

```
g++ -v
```

 Important

Some newer library features are statically linked into applications built with Red Hat Developer Toolset to support execution on multiple versions of Red Hat Enterprise Linux. This adds a small additional security risk as normal Red Hat Enterprise Linux errata would not change this code. If the need for developers to rebuild their applications due to such an issue arises, Red Hat will signal this via a security erratum. Developers are strongly advised not to statically link their entire application for the same reasons.

### Example 2.3. Compiling a C++ Program on the Command Line

Imagine that you have a source file named **hello.cpp** with the following contents:

```
#include <iostream>

using namespace std;

int main(int argc, char *argv[]) {
    cout << "Hello, World!" << endl;
    return 0;
}
```

To compile this source code on the command line by using the **g++** compiler from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-1.1 'g++ -o hello hello.cpp'
```

This creates a new binary file called **hello** in the current working directory.

### 2.2.3. Running a C++ Program

When **g++** compiles a program, it creates an executable binary file. To run this program on the command line, change to the directory with the executable file and type:

```
./file_name
```

### Example 2.4. Running a C++ Program on the Command Line

Assuming that you have successfully compiled the `hello` binary file as shown in [Example 2.3, “Compiling a C++ Program on the Command Line”](#), you can run it by typing the following at a shell prompt:

```
~]$ ./hello
Hello, World!
```

## 2.3. GNU Fortran Compiler

### 2.3.1. Installing the Fortran Compiler

In Red Hat Developer Toolset, the GNU Fortran compiler is provided by the `devtoolset-1.1-gcc-gfortran` package, and is automatically installed with `devtoolset-1.1` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 2.3.2. Using the Fortran Compiler

To compile a Fortran program on the command line, run the `gfortran` compiler as follows:

```
sc1 enable devtoolset-1.1 'gfortran -o output_file source_file...'
```

This creates a binary file named `output_file` in the current working directory. If the `-o` option is omitted, the compiler creates a file named `a.out` by default.

When you are working on a project that consists of several source files, it is common to compile an object file for each of the source files first and then link these object files together. This way, when you change a single source file, you can recompile only this file without having to compile the entire project. To compile an object file on the command line, run the following command:

```
sc1 enable devtoolset-1.1 'gfortran -o object_file -c source_file'
```

This creates an object file named `object_file`. If the `-o` option is omitted, the compiler creates a file named after the source file with the `.o` file extension. To link object files together and create a binary file, run:

```
sc1 enable devtoolset-1.1 'gfortran -o output_file object_file...'
```

Note that you can execute any command using the `sc1` utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset `gfortran` as default:

```
sc1 enable devtoolset-1.1 'bash'
```





## Note

To verify the version of **gfortran** you are using at any point, type the following at a shell prompt:

```
which gfortran
```

Red Hat Developer Toolset's **gfortran** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **gfortran**:

```
gfortran -v
```



## Important

Some newer library features are statically linked into applications built with Red Hat Developer Toolset to support execution on multiple versions of Red Hat Enterprise Linux. This adds a small additional security risk as normal Red Hat Enterprise Linux errata would not change this code. If the need for developers to rebuild their applications due to such an issue arises, Red Hat will signal this via a security erratum. Developers are strongly advised not to statically link their entire application for the same reasons.

### Example 2.5. Compiling a Fortran Program on the Command Line

Imagine that you have a source file named **hello.f** with the following contents:

```
program hello
  print *, "Hello, World!"
end program hello
```

To compile this source code on the command line by using the **gfortran** compiler from Red Hat Developer Toolset, type:

```
~]$ scl enable devtoolset-1.1 'gfortran -o hello hello.f'
```

This creates a new binary file called **hello** in the current working directory.

### 2.3.3. Running a Fortran Program

When **gfortran** compiles a program, it creates an executable binary file. To run this program on the command line, change to the directory with the executable file and type:

```
./file_name
```

### Example 2.6. Running a Fortran Program on the Command Line

Assuming that you have successfully compiled the `hello` binary file as shown in [Example 2.5, “Compiling a Fortran Program on the Command Line”](#), you can run it by typing the following at a shell prompt:

```
~]$ ./hello
Hello, World!
```

## 2.4. Additional Resources

A detailed description of the GNU Compiler Collections and its features is beyond the scope of this book. For more information, refer to the resources listed below.

### Installed Documentation

- ▶ **gcc(1)** — The manual page for the **gcc** compiler provides detailed information on its usage; with few exceptions, **g++** accepts the same command line options as **gcc**. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-1.1 'man gcc'
```

- ▶ **gfortran(1)** — The manual page for the **gfortran** compiler provides detailed information on its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-1.1 'man gfortran'
```

- ▶ *C++ Standard Library Documentation* — Documentation on the C++ standard library can be optionally installed by typing the following at a shell prompt as **root**:

```
yum install devtoolset-1.1-libstdc++-docs
```

Once installed, HTML documentation is available at `/opt/rh/devtoolset-1.1/root/usr/share/doc/devtoolset-1.1-libstdc++-docs-4.7.2/html/index.html`.

### Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides in-depth information about GCC.
- ▶ [Using the GNU Compiler Collection](#) — The official GCC manual provides an in-depth description of the GNU compilers and their usage.
- ▶ [The GNU C++ Library](#) — The GNU C++ library documentation provides detailed information about the GNU implementation of the standard C++ library.
- ▶ [The GNU Fortran Compiler](#) — The GNU Fortran compiler documentation provides detailed information on **gfortran**'s usage.

### See Also

- ▶ [Section B.1 “Changes in GCC 4.7.2”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of the GNU Compiler Collection, as well as information about the language, ABI, and debugging compatibility.

- ▶ [Chapter 1, \*Red Hat Developer Toolset\*](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 3, \*GNU Debugger \(GDB\)\*](#) provides information on how to debug programs written in C, C++, and Fortran.
- ▶ [Chapter 4, \*binutils\*](#) explains how to use the binutils, a collection of binary tools to inspect and manipulate object files and binaries.
- ▶ [Chapter 5, \*elfutils\*](#) explains how to use elfutils, a collection of binary tools to inspect and manipulate ELF files.
- ▶ [Chapter 10, \*Eclipse\*](#) provides a general introduction to the Eclipse development environment, and describes how to use it with the tools from Red Hat Developer Toolset.

## Chapter 3. GNU Debugger (GDB)

The **GNU Debugger**, commonly abbreviated as **GDB**, is a command line tool that can be used to debug programs written in various programming languages. It allows you to inspect memory within the code being debugged, control the execution state of the code, detect the execution of particular sections of code, and much more.

Red Hat Developer Toolset is distributed with **GDB 7.5**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements, including improved support for Python scripting, ambiguous line specifications, and tracepoints, as well as improved inferior control commands and handling of C++ debuggee executables. For a detailed list of changes, refer to [Section B.2, “Changes in GDB 7.5”](#).

### 3.1. Installing the GNU Debugger

In Red Hat Developer Toolset, the GNU Debugger is provided by the *devtoolset-1.1-gdb* package, and is automatically installed with *devtoolset-1.1* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 3.2. Preparing a Program for Debugging

#### Compiling Programs with Debugging Information

To compile a C program with debugging information that can be read by the GNU Debugger, make sure the **gcc** compiler is run with the **-g** option. To do so on the command line, use a command in the following form:

```
scl enable devtoolset-1.1 'gcc -g -o output_file input_file...'
```

Similarly, to compile a C++ program with debugging information, run:

```
scl enable devtoolset-1.1 'g++ -g -o output_file input_file...'
```

### Example 3.1. Compiling a C Program With Debugging Information

Imagine you have a source file named **fibonacci.c** with the following contents:

```
#include <stdio.h>
#include <limits.h>

int main (int argc, char *argv[]) {
    unsigned long int a = 0;
    unsigned long int b = 1;
    unsigned long int sum;

    while (b < LONG_MAX) {
        printf("%ld ", b);
        sum = a + b;
        a = b;
        b = sum;
    }

    return 0;
}
```

To compile this program on the command line using GCC from Red Hat Developer Toolset with debugging information for the GNU Debugger, type:

```
~]$ scl enable devtoolset-1.1 'gcc -g -o fibonacci fibonacci.c'
```

This creates a new binary file called **fibonacci** in the current working directory.

### Installing Debugging Information for Existing Packages

To install debugging information for a package that is already installed on the system, type the following at a shell prompt as **root**:

```
debuginfo-install package_name
```

Note that the *yum-utils* package must be installed for the **debuginfo-install** utility to be available on your system.

### Example 3.2. Installing Debugging Information for the *glibc* Package

To install debugging information for the *glibc* package, type:

```
~]# debuginfo-install glibc
Loaded plugins: product-id, refresh-packagekit, subscription-manager
--> Running transaction check
---> Package glibc-debuginfo.x86_64 0:2.12-1.47.el6_2.5 will be installed
...
```

## 3.3. Running the GNU Debugger

To run the GNU Debugger on a program you want to debug, type the following at a shell prompt:

```
sc1 enable devtoolset-1.1 'gdb file_name'
```

This starts the **gdb** debugger in interactive mode and displays the default prompt, **(gdb)**. To quit the debugging session and return to the shell prompt, run the following command at any time:

```
quit
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **gdb** as default:

```
sc1 enable devtoolset-1.1 'bash'
```



### Note

To verify the version of **gdb** you are using at any point, type the following at a shell prompt:

```
which gdb
```

Red Hat Developer Toolset's **gdb** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **gdb**:

```
gdb -v
```

### Example 3.3. Running the gdb Utility on the fibonacci Binary File

Assuming that you have successfully compiled the **fibonacci** binary file as shown in [Example 3.1, "Compiling a C Program With Debugging Information"](#), you can start debugging it with **gdb** by typing the following at a shell prompt:

```
~]$ sc1 enable devtoolset-1.1 'gdb fibonacci'
GNU gdb (GDB) Red Hat Enterprise Linux (7.4.50.20120120-43.el6)
Copyright (C) 2012 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law. Type "show copying"
and "show warranty" for details.
This GDB was configured as "x86_64-redhat-linux-gnu".
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>.
(gdb)
```

## 3.4. Listing Source Code

To view the source code of the program you are debugging, run the following command:

**list**

Before you start the execution of the program you are debugging, **gdb** displays first ten lines of the source code and any subsequent use of this command lists another ten lines. Once you start the execution, **gdb** displays the lines that are surrounding the line on which the execution stops, typically when you set a breakpoint.

You can also display the code that is surrounding a particular line. To do so, run the command in the following form:

```
list [file_name:]line_number
```

Similarly, to display the code that is surrounding the beginning of a particular function, run:

```
list [file_name:]function_name
```

Note that you can change the number of lines the **list** command displays by running the following command:

```
set listsize number
```

**Example 3.4. Listing the Source Code of the fibonacci Binary File**

The **fibonacci.c** file listed in [Example 3.1, “Compiling a C Program With Debugging Information”](#) has exactly 17 lines. Assuming that you have compiled it with debugging information and you want the **gdb** utility to be capable of listing the entire source code, you can run the following command to change the number of listed lines to 20:

```
(gdb) set listsize 20
```

You can now display the entire source code of the file you are debugging by running the **list** command with no additional arguments:

```
(gdb) list
1      #include <stdio.h>
2      #include <limits.h>
3
4      int main (int argc, char *argv[]) {
5          unsigned long int a = 0;
6          unsigned long int b = 1;
7          unsigned long int sum;
8
9          while (b < LONG_MAX) {
10             printf("%ld ", b);
11             sum = a + b;
12             a = b;
13             b = sum;
14         }
15
16         return 0;
17     }
```

## Setting a New Breakpoint

To set a new breakpoint at a certain line, run the following command:

```
break [file_name:]line_number
```

You can also set a breakpoint on a certain function:

```
break [file_name:]function_name
```

### Example 3.5. Setting a New Breakpoint

Assuming that you have compiled the `fibonacci.c` file listed in [Example 3.1, “Compiling a C Program With Debugging Information”](#) with debugging information, you can set a new breakpoint at line 10 by running the following command:

```
(gdb) break 10
Breakpoint 1 at 0x4004e5: file fibonacci.c, line 10.
```

## Listing Breakpoints

To display a list of currently set breakpoints, run the following command:

```
info breakpoints
```

### Example 3.6. Listing Breakpoints

Assuming that you have followed the instructions in [Example 3.5, “Setting a New Breakpoint”](#), you can display the list of currently set breakpoints by running the following command:

```
(gdb) info breakpoints
Num      Type           Disp Enb Address                What
1        breakpoint     keep y   0x00000000004004e5  in main at fibonacci.c:10
```

## Deleting Existing Breakpoints

To delete a breakpoint that is set at a certain line, run the following command:

```
clear line_number
```

Similarly, to delete a breakpoint that is set on a certain function, run:

```
clear function_name
```

You can also delete all breakpoints at once. To do so, run the `clear` command with no additional arguments:

```
clear
```



### Example 3.7. Deleting an Existing Breakpoint

Imagine that you have accidentally set a breakpoint at a wrong line and you want to remove it. Assuming that you have compiled the `fibonacci.c` file listed in [Example 3.1, “Compiling a C Program With Debugging Information”](#) with debugging information, you can set a new breakpoint at line 7 by running the following command:

```
(gdb) break 7
Breakpoint 2 at 0x4004e3: file fibonacci.c, line 7.
```

To remove this breakpoint, type:

```
(gdb) clear 7
Deleted breakpoint 2
```

## 3.6. Starting Execution

To start execution of the program you are debugging, run the following command:

```
run
```

If the program accepts any command line arguments, you can provide them as arguments to the `run` command:

```
run argument...
```

The execution stops when a first breakpoint (if any) is reached, when an error occurs, or when the program terminates.

### Example 3.8. Executing the fibonacci Binary File

Assuming that you have followed the instructions in [Example 3.5, “Setting a New Breakpoint”](#), you can execute the `fibonacci` binary file by running the following command:

```
(gdb) run
Starting program: /home/john/fibonacci

Breakpoint 1, main (argc=1, argv=0x7fffffffef4d8) at fibonacci.c:10
10     printf("%ld ", b);
```

## 3.7. Displaying Current Values

The `gdb` utility allows you to display the value of almost anything that is relevant to the program, from a variable of any complexity to a valid expression or even a library function. However, the most common task is to display the value of a variable.

To display the current value of a certain variable, run the following command:

```
print variable_name
```

### Example 3.9. Displaying the Current Values of Variables

Assuming that you have followed the instructions in [Example 3.8, “Executing the fibonacci Binary File”](#) and the execution of the `fibonacci` binary stopped after reaching the breakpoint at line 10, you can display the current values of variables `a` and `b` as follows:

```
(gdb) print a  
$1 = 0  
(gdb) print b  
$2 = 1
```

## 3.8. Continuing Execution

To resume the execution of the program you are debugging after it reached a breakpoint, run the following command:

```
continue
```

The execution stops again when another breakpoint is reached. To skip a certain number of breakpoints (typically when you are debugging a loop), you can run the `continue` command in the following form:

```
continue number
```

The `gdb` utility also allows you to stop the execution after executing a single line of code. To do so, run:

```
step
```

Finally, you can execute a certain number of lines by using the `step` command in the following form:

```
step number
```

### Example 3.10. Continuing the Execution of the fibonacci Binary File

Assuming that you have followed the instructions in [Example 3.8, “Executing the fibonacci Binary File”](#) and the execution of the `fibonacci` binary stopped after reaching the breakpoint at line 10, you can resume the execution by running the following command:

```
(gdb) continue
Continuing.

Breakpoint 1, main (argc=1, argv=0x7fffffff4d8) at fibonacci.c:10
10         printf("%ld ", b);
```

The execution stops the next time the breakpoint is reached. To execute next three lines of code, type:

```
(gdb) step 3
13         b = sum;
```

This allows you to verify the current value of the `sum` variable before it is assigned to `b`:

```
(gdb) print sum
$3 = 2
```

## 3.9. Additional Resources

A detailed description of the GNU Debugger and all its features is beyond the scope of this book. For more information, refer to the resources listed below.

### Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on the GNU Debugger and debugging.
- ▶ [GDB Documentation](#) — The official GDB documentation includes the *GDB User Manual* and other reference material.

### See Also

- ▶ [Section B.2, “Changes in GDB 7.5”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of the GNU Debugger.
- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 2, GNU Compiler Collection \(GCC\)](#) provides further information on how to compile programs written in C, C++, and Fortran.
- ▶ [Chapter 10, Eclipse](#) provides a general introduction to the Eclipse development environment, and describes how to use it with the tools from Red Hat Developer Toolset.

## Chapter 4. binutils

**binutils** is a collection of various binary tools such as the GNU linker, GNU assembler, and other utilities that allow you to inspect and manipulate object files and binaries. Refer to [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#) for a complete list of binary tools that are distributed with the Red Hat Developer Toolset version of binutils.

Red Hat Developer Toolset is distributed with **binutils 2.23.51**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements, including the new **gold** linker, several new command line options, improvements to the linker script language, and support for link-time optimization, compressed debug sections, and new instruction sets. For a detailed list of changes, refer to [Section B.3, “Changes in binutils 2.23.51”](#).

**Table 4.1. Tools Included in binutils for Red Hat Developer Toolset**

Name	Description
<b>addr2line</b>	Translates addresses into file names and line numbers.
<b>ar</b>	Creates, modifies, and extracts files from archives.
<b>as</b>	The GNU assembler.
<b>c++filt</b>	Decodes mangled C++ symbols.
<b>elfedit</b>	Examines and edits ELF files.
<b>gprof</b>	Display profiling information.
<b>ld</b>	The GNU linker.
<b>ld.bfd</b>	An alternative to the GNU linker.
<b>ld.gold</b>	A new ELF linker.
<b>nm</b>	Lists symbols from object files.
<b>objcopy</b>	Copies and translates object files.
<b>objdump</b>	Displays information from object files.
<b>ranlib</b>	Generates an index to the contents of an archive to make access to this archive faster.
<b>readelf</b>	Displays information about ELF files.
<b>size</b>	Lists section sizes of object or archive files.
<b>strings</b>	Displays printable character sequences in files.
<b>strip</b>	Discards all symbols from object files.

### 4.1. Installing binutils

In Red Hat Developer Toolset, binutils are provided by the *devtoolset-1.1-binutils* package, and are automatically installed with *devtoolset-1.1* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 4.2. Using the GNU Assembler

To produce an object file from an assembly language program, run the **as** tool as follows:

```
scl enable devtoolset-1.1 'as [option...] -o object_file source_file'
```

This creates an object file named ***object\_file*** in the current working directory.

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **as** as default:

```
scl enable devtoolset-1.1 'bash'
```

### Note

To verify the version of **as** you are using at any point, type the following at a shell prompt:

```
which as
```

Red Hat Developer Toolset's **as** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **as**:

```
as -v
```

## 4.3. Using the GNU Linker

To create an executable binary file or a library from object files, run the **ld** tool as follows:

```
scl enable devtoolset-1.1 'ld [option...] -o output_file object_file...'
```

This creates a binary file named **output\_file** in the current working directory. If the **-o** option is omitted, the compiler creates a file named **a.out** by default.

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **ld** as default:

```
scl enable devtoolset-1.1 'bash'
```

### Note

To verify the version of **ld** you are using at any point, type the following at a shell prompt:

```
which ld
```

Red Hat Developer Toolset's **ld** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **ld**:

```
ld -v
```

## 4.4. Using Other Binary Tools

The binutils provide many binary tools other than a linker and assembler. For a complete list of these

tools, refer to [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#).

To execute any of the tools that are part of binutils, run the command as follows:

```
scl enable devtoolset-1.1 'tool [option...] file_name'
```

Refer to [Table 4.1, “Tools Included in binutils for Red Hat Developer Toolset”](#) for a list of tools that are distributed with binutils. For example, to use the **objdump** tool to inspect an object file, type:

```
scl enable devtoolset-1.1 'objdump [option...] object_file'
```

Note that you can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset binary tools as default:

```
scl enable devtoolset-1.1 'bash'
```



### Note

To verify the version of binutils you are using at any point, type the following at a shell prompt:

```
which objdump
```

Red Hat Developer Toolset's **objdump** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **objdump**:

```
objdump -v
```

## 4.5. Additional Resources

A detailed description of binutils is beyond the scope of this book. For more information, refer to the resources listed below.

### Installed Documentation

- **as(1)**, **ld(1)**, **addr2line(1)**, **ar(1)**, **c++filt(1)**, **elfedit(1)**, **gprof(1)**, **nm(1)**, **objcopy(1)**, **objdump(1)**, **ranlib(1)**, **readelf(1)**, **size(1)**, **strings(1)**, **strip(1)**, — Manual pages for various binutils tools provide more information about their respective usage. To display a manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-1.1 'man tool'
```

### Online Documentation

- [Documentation for binutils](#) — The binutils documentation provides an in-depth description of the binary tools and their usage.

### See Also

- [Section B.3 “Changes in binutils 2.23.51”](#) provides a comprehensive list of features and improvements over the Red Hat Enterprise Linux system version of binutils.

- ▶ [Chapter 1, \*Red Hat Developer Toolset\*](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 5, \*elfutils\*](#) explains how to use elfutils, a collection of binary tools to inspect and manipulate ELF files.
- ▶ [Chapter 2, \*GNU Compiler Collection \(GCC\)\*](#) provides information on how to compile programs written in C, C++, and Fortran.

## Chapter 5. elfutils

**elfutils** is a collection of various binary tools such as **eu-objdump**, **eu-readelf**, and other utilities that allow you to inspect and manipulate ELF files. Refer to [Table 5.1, “Tools Included in elfutils for Red Hat Developer Toolset”](#) for a complete list of binary tools that are distributed with the Red Hat Developer Toolset version of elfutils.

Red Hat Developer Toolset is distributed with **elfutils 0.154**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements.

**Table 5.1. Tools Included in elfutils for Red Hat Developer Toolset**

Name	Description
<b>eu-addr2line</b>	Translates addresses into file names and line numbers.
<b>eu-ar</b>	Creates, modifies, and extracts files from archives.
<b>eu-elfcmp</b>	Compares relevant parts of two ELF files for equality.
<b>eu-elflint</b>	Verifies that ELF files are compliant with the <i>generic ABI</i> (gABI) and <i>processor-specific supplement ABI</i> (psABI) specification.
<b>eu-findtextrel</b>	Locates the source of text relocations in files.
<b>eu-make-debug-archive</b>	Creates an offline archive for debugging.
<b>eu-nm</b>	Lists symbols from object files.
<b>eu-objdump</b>	Displays information from object files.
<b>eu-ranlib</b>	Generates an index to the contents of an archive to make access to this archive faster.
<b>eu-readelf</b>	Displays information about ELF files.
<b>eu-size</b>	Lists section sizes of object or archive files.
<b>eu-strings</b>	Displays printable character sequences in files.
<b>eu-strip</b>	Discards all symbols from object files.
<b>eu-unstrip</b>	Combines stripped files with separate symbols and debug information.

### 5.1. Installing elfutils

In Red Hat Developer Toolset, elfutils is provided by the *devtoolset-1.1-elfutils* package, and is automatically installed with *devtoolset-1.1* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 5.2. Using elfutils

To execute any of the tools that are part of elfutils, run the command as follows:

```
sc1 enable devtoolset-1.1 'tool [option...] file_name'
```

Refer to [Table 5.1, “Tools Included in elfutils for Red Hat Developer Toolset”](#) for a list of tools that are distributed with elfutils. For example, to use the **eu-objdump** tool to inspect an object file, type:

```
sc1 enable devtoolset-1.1 'eu-objdump [option...] object_file'
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat



Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset binary tools as default:

```
scl enable devtoolset-1.1 'bash'
```



## Note

To verify the version of elfutils you are using at any point, type the following at a shell prompt:

```
which eu-objdump
```

Red Hat Developer Toolset's **eu-objdump** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **eu-objdump**:

```
eu-objdump -V
```

## 5.3. Additional Resources

A detailed description of elfutils is beyond the scope of this book. For more information, refer to the resources listed below.

### See Also

- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 4, binutils](#) explains how to use the binutils, a collection of binary tools to inspect and manipulate object files and binaries.
- ▶ [Chapter 2, GNU Compiler Collection \(GCC\)](#) provides information on how to compile programs written in C, C++, and Fortran.

## Chapter 6. dwz

**dwz** is a command line tool that attempts to optimize DWARF debugging information contained in ELF shared libraries and ELF executables for size. To do so, **dwz** replaces DWARF information representation with equivalent smaller representation where possible, and reduces the amount of duplication by using techniques from *Appendix E* of the *DWARF Standard*.

Red Hat Developer Toolset is distributed with **dwz 0.7**.

### 6.1. Installing dwz

In Red Hat Developer Toolset, the **dwz** utility is provided by the `devtoolset-1.1-dwz` package, and is automatically installed with `devtoolset-1.1` as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 6.2. Using dwz

To optimize DWARF debugging information in a binary file, run the **dwz** tool as follows:

```
sc1 enable devtoolset-1.1 'dwz [option...] file_name'
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset **dwz** as default:

```
sc1 enable devtoolset-1.1 'bash'
```



#### Note

To verify the version of **dwz** you are using at any point, type the following at a shell prompt:

```
which dwz
```

Red Hat Developer Toolset's **dwz** executable path will begin with `/opt`. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **dwz**:

```
dwz -v
```

### 6.3. Additional Resources

A detailed description of **dwz** and its features is beyond the scope of this book. For more information, refer to the resources listed below.

#### Installed Documentation

- **dwz(1)** — The manual page for the **dwz** utility provides detailed information on its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-1.1 'man dwz'
```

### See Also

- ▶ [Chapter 1, \*Red Hat Developer Toolset\*](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 2, \*GNU Compiler Collection \(GCC\)\*](#) provides information on how to compile programs written in C, C++, and Fortran.
- ▶ [Chapter 4, \*binutils\*](#) explains how to use the binutils, a collection of binary tools to inspect and manipulate object files and binaries.
- ▶ [Chapter 5, \*elfutils\*](#) explains how to use elfutils, a collection of binary tools to inspect and manipulate ELF files.

## Part III. Performance Monitoring Tools

## Chapter 7. SystemTap

**SystemTap** is a tracing and probing tool that allows users to monitor the activities of the entire system without needing to instrument, recompile, install, and reboot. It is programmable with a custom scripting language, which gives it expressiveness (to trace, filter, and analyze) and reach (to look into the running kernel and applications).

SystemTap can monitor various types of events, such as function calls within the kernel or applications, timers, tracepoints, performance counters, and so on. Some included example scripts produce output similar to **netstat**, **ps**, **top**, and **iostat**, others include pretty-printed function callgraph traces or tools for working around security bugs.

Red Hat Developer Toolset is distributed with **SystemTap 1.8**.

**Table 7.1. Tools Distributed with SystemTap for Red Hat Developer Toolset**

Name	Description
<b>stap</b>	Translates probing instructions into C code, builds a kernel module, and loads it into a running Linux kernel.
<b>staprun</b>	Loads, unloads, attaches to, and detaches from kernel modules built with the <b>stap</b> utility.
<b>stapsh</b>	Serves as a remote shell for SystemTap.
<b>stap-prep</b>	Determines and—if possible—downloads the kernel information packages that are required to run SystemTap.
<b>stap-merge</b>	Merges per-CPU files. This script is automatically executed when the <b>stap</b> utility is executed with the <b>-b</b> command line option.
<b>stap-report</b>	Gathers important information about the system for the purpose of reporting a bug in SystemTap.

### 7.1. Installing SystemTap

In Red Hat Developer Toolset, **SystemTap** is provided by the *devtoolset-1.1-systemtap* package, and is automatically installed with *devtoolset-1.1* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

In order to place instrumentation into the Linux kernel, SystemTap may also require installation of additional packages with debugging information. To determine which packages to install, run the **stap-prep** utility as follows:

```
scl enable devtoolset-1.1 'stap-prep'
```

Note that if you execute this command as the **root** user, the utility automatically offers the packages for installation. For more information on how to install these packages on your system, refer to the *Red Hat Enterprise Linux SystemTap Beginners Guide*.

### 7.2. Using SystemTap

To execute any of the tools that are part of SystemTap, type the following at a shell prompt:

```
scl enable devtoolset-1.1 'tool [option...]
```

Refer to [Table 7.1, “Tools Distributed with SystemTap for Red Hat Developer Toolset”](#) for a list of tools that are distributed with SystemTap. For example, to run the **stap** tool to build an instrumentation module, type:

```
sc1 enable devtoolset-1.1 'stap [option...] argument...'
```

Note that you can execute any command using the **sc1** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset SystemTap as default:

```
sc1 enable devtoolset-1.1 'bash'
```



### Note

To verify the version of SystemTap you are using at any point, type the following at a shell prompt:

```
which stap
```

Red Hat Developer Toolset's **stap** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset SystemTap:

```
stap -V
```

## 7.3. Additional Resources

A detailed description of SystemTap and its features is beyond the scope of this book. For more information, refer to the resources listed below.

### Installed Documentation

- ▶ **stap(1)** — The manual page for the **stap** command provides detailed information on its usage, as well as references to other related manual pages. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
sc1 enable devtoolset-1.1 'man stap'
```

- ▶ **staprun(8)** — The manual page for the **staprun** command provides detailed information on its usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
sc1 enable devtoolset-1.1 'man staprun'
```

- ▶ *SystemTap Tapset Reference Manual* — HTML documentation on the most common tapset definitions is located at `/opt/rh/devtoolset-1.1/root/usr/share/doc/devtoolset-1.1-systemtap-client-1.8/index.html`.

### Online Documentation

- ▶ [Red Hat Enterprise Linux 6 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for Red Hat Enterprise Linux 6 provides an introduction to SystemTap and its usage.
- ▶ [Red Hat Enterprise Linux 5 SystemTap Beginners Guide](#) — The *SystemTap Beginners Guide* for

Red Hat Enterprise Linux 5 provides an introduction to SystemTap and its usage.

- ▶ [Red Hat Enterprise Linux 6 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 6 provides further details about SystemTap.
- ▶ [Red Hat Enterprise Linux 5 SystemTap Tapset Reference](#) — The *SystemTap Tapset Reference* for Red Hat Enterprise Linux 5 provides further details about SystemTap.
- ▶ [The SystemTap Documentation](#) — The official SystemTap documentation provides further documentation on SystemTap, as well as numerous examples of SystemTap scripts.

### See Also

- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 8, Valgrind](#) explains how to use Valgrind to profile applications and detect memory errors and memory management problems, such as the use of uninitialized memory, improper allocation and freeing of memory, and the use of improper arguments in system calls.
- ▶ [Chapter 9, OProfile](#) explains how to use OProfile to determine which sections of code consume the greatest amount of CPU time and why.

## Chapter 8. Valgrind

**Valgrind** is an instrumentation framework that ships with a number of tools to profile applications. It can be used to detect various memory errors and memory management problems, such as the use of uninitialized memory or improper allocation and freeing of memory, or to identify the use of improper arguments in system calls. For a complete list of profiling tools that are distributed with the Red Hat Developer Toolset version of Valgrind, refer to [Table 8.1, “Tools Distributed with Valgrind for Red Hat Developer Toolset”](#).

Valgrind profiles an application by rewriting it and instrumenting the rewritten binary. This allows you to profile your application without the need to recompile it, but it also makes Valgrind significantly slower than other profilers, especially when performing extremely detailed runs. It is therefore not suited to debugging time-specific issues, or kernel-space debugging.

Red Hat Developer Toolset is distributed with **Valgrind 3.8.1**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements.

**Table 8.1. Tools Distributed with Valgrind for Red Hat Developer Toolset**

Name	Description
<b>Memcheck</b>	Detects memory management problems by intercepting system calls and checking all read and write operations.
<b>Cachegrind</b>	Identifies the sources of cache misses by simulating the level 1 instruction cache (I1), level 1 data cache (D1), and unified level 2 cache (L2).
<b>Callgrind</b>	Generates a call graph representing the function call history.
<b>Helgrind</b>	Detects synchronization errors in multithreaded C, C++, and Fortran programs that use POSIX threading primitives.
<b>DRD</b>	Detects errors in multithreaded C and C++ programs that use POSIX threading primitives or any other threading concepts that are built on top of these POSIX threading primitives.
<b>Massif</b>	Monitors heap and stack usage.

### 8.1. Installing Valgrind

In Red Hat Developer Toolset, Valgrind is provided by the *devtoolset-1.1-valgrind* package, and is automatically installed with *devtoolset-1.1*. If you intend to use Valgrind to profile parallel programs that use the Message Passing Interface (MPI) protocol, install the *devtoolset-1.1-valgrind-openmpi* package.

For detailed information about how to install Red Hat Developer Toolset and related packages to your system, refer to [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 8.2. Using Valgrind

To run any of the Valgrind tools on a program you want to profile, type the following at a shell prompt:

```
scl enable devtoolset-1.1 'valgrind [--tool=tool] program [option...]
```

Refer to [Table 8.1, “Tools Distributed with Valgrind for Red Hat Developer Toolset”](#) for a list of tools that are distributed with Valgrind. The argument of the `--tool` command line option must be specified in lower case, and if this option is omitted, Valgrind uses **Memcheck** by default. For example, to run Cachegrind on a program to identify the sources of cache misses, type:



```
scl enable devtoolset-1.1 'valgrind --tool=cachegrind program [option...]'
```

Note that you can execute any command using the `scl` utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset Valgrind as default:

```
scl enable devtoolset-1.1 'bash'
```



### Note

To verify the version of Valgrind you are using at any point, type the following at a shell prompt:

```
which valgrind
```

Red Hat Developer Toolset's **valgrind** executable path will begin with `/opt`. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset Valgrind:

```
valgrind --version
```

## 8.3. Additional Resources

A detailed description of Valgrind and its features is beyond the scope of this book. For more information, refer to the resources listed below.

### Installed Documentation

- ▶ **valgrind(1)** — The manual page for the **valgrind** utility provides detailed information on how to use Valgrind. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
scl enable devtoolset-1.1 'man valgrind'
```

- ▶ *Valgrind Documentation* — HTML documentation for Valgrind is located at `/opt/rh/devtoolset-1.1/root/usr/share/doc/devtoolset-1.1-valgrind-3.8.1/html/index.html`.

### Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information about Valgrind and its Eclipse plug-in.
- ▶ [Red Hat Enterprise Linux 6 Performance Tuning Guide](#) — The *Performance Tuning Guide* for Red Hat Enterprise Linux 6 provides more detailed information about using Valgrind to profile applications.

### See Also

- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 7, SystemTap](#) provides an introduction to SystemTap and explains how to use it to monitor the activities of a running system.
- ▶ [Chapter 9, OProfile](#) explains how to use OProfile to determine which sections of code consume the greatest amount of CPU time and why.



## Chapter 9. OProfile

**OProfile** is a low overhead, system-wide profiler that 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 level 2 cache (L2) requests, and the number of hardware interrupts received. It consists of a configuration utility, a daemon for collecting data, and a number of tools that can be used to transform the data into a human-readable form. For a complete list of tools that are distributed with the Red Hat Developer Toolset version of OProfile, refer to [Table 9.1, “Tools Distributed with OProfile for Red Hat Developer Toolset”](#).

OProfile profiles an application without adding any instrumentation by recording the details of every nth event. This allows it to consume fewer resources than Valgrind, but also causes its samples to be less precise. Unlike Valgrind, which only collects data for a single process and its children in user-space, OProfile is well suited to collect system-wide data on both user-space and kernel-space processes, and requires **root** privileges to run.

Red Hat Developer Toolset is distributed with **OProfile 0.9.7**. This version is more recent than the version included in Red Hat Enterprise Linux and provides numerous bug fixes and enhancements.

**Table 9.1. Tools Distributed with OProfile for Red Hat Developer Toolset**

Name	Description
<b>oprofiled</b>	The OProfile daemon that collects profiling data.
<b>opcontrol</b>	Starts, stops, and configures the OProfile daemon.
<b>opannotate</b>	Generates an annotated source file or assembly listing from the profiling data.
<b>oparchive</b>	Generates a directory containing executable, debug, and sample files.
<b>opgprof</b>	Generates a summary of a profiling session in a format compatible with <b>gprof</b> .
<b>ophelp</b>	Displays a list of available events.
<b>opimport</b>	Converts a sample database file from a foreign binary format to the native format.
<b>opjitconv</b>	Converts a just-in-time (JIT) dump file to the Executable and Linkable Format (ELF).
<b>opreport</b>	Generates image and symbol summaries of a profiling session.

### 9.1. Installing OProfile

In Red Hat Developer Toolset, OProfile is provided by the *devtoolset-1.1-oprofile* package, and is automatically installed with *devtoolset-1.1* as described in [Section 1.5, “Installing Red Hat Developer Toolset”](#).

### 9.2. Using OProfile

To run any of the tools that are distributed with OProfile, type the following at a shell prompt as **root**:

```
sc1 enable devtoolset-1.1 'tool [option...]
```

Refer to [Table 9.1, “Tools Distributed with OProfile for Red Hat Developer Toolset”](#) for a list of tools that are distributed with Valgrind. For example, to use the **ophelp** command to list available events in the XML format, type:

```
sc1 enable devtoolset-1.1 'ophelp -X'
```

Note that you can execute any command using the `sc1` utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalent. This allows you to run a shell session with Red Hat Developer Toolset OProfile as default:

```
sc1 enable devtoolset-1.1 'bash'
```



### Note

To verify the version of OProfile you are using at any point, type the following at a shell prompt:

```
which opcontrol
```

Red Hat Developer Toolset's `opcontrol` executable path will begin with `/opt`. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset OProfile:

```
opcontrol --version
```

## 9.3. Additional Resources

A detailed description of OProfile and its features is beyond the scope of this book. For more information, refer to the resources listed below.

### Installed Documentation

- ▶ `oprofile(1)` — The manual page named `oprofile` provides an overview of OProfile and available tools. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
sc1 enable devtoolset-1.1 'man oprofile'
```

- ▶ `opannotate(1)`, `oparchive(1)`, `opcontrol(1)`, `opgprof(1)`, `ophelp(1)`, `opimport(1)`, `opreport(1)` — Manual pages for various tools distributed with OProfile provide more information on their respective usage. To display the manual page for the version included in Red Hat Developer Toolset, type:

```
sc1 enable devtoolset-1.1 'man tool'
```

### Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on OProfile.
- ▶ [Red Hat Enterprise Linux 6 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 6 describes in detail how to install, configure, and start using OProfile on this system.
- ▶ [Red Hat Enterprise Linux 5 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 5 describes in detail how to install, configure, and start using OProfile on this system.

### See Also

- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.

- ▶ [Chapter 7, \*SystemTap\*](#) provides an introduction to SystemTap and explains how to use it to monitor the activities of a running system.
- ▶ [Chapter 8, \*Valgrind\*](#) explains how to use Valgrind to profile applications and detect memory errors and memory management problems, such as the use of uninitialized memory, improper allocation and freeing of memory, and the use of improper arguments in system calls.

## Part IV. Integrated Development Environments

## Chapter 10. Eclipse

**Eclipse** is a powerful development environment that provides tools for each phase of the development process. It integrates a variety of disparate tools into a unified environment to create a rich development experience, provides a fully configurable user interface, and features a pluggable architecture that allows for extension in a variety of ways. For instance, the Valgrind plug-in allows programmers to perform memory profiling, otherwise performed on the command line, through the Eclipse user interface.

Eclipse provides a graphical development environment alternative to traditional interaction with command line tools and as such, it is a welcome alternative to developers who do not want to use the command line interface. The traditional, mostly command line based Linux tools suite (such as **gcc** or **gdb**) and Eclipse offer two distinct approaches to programming.

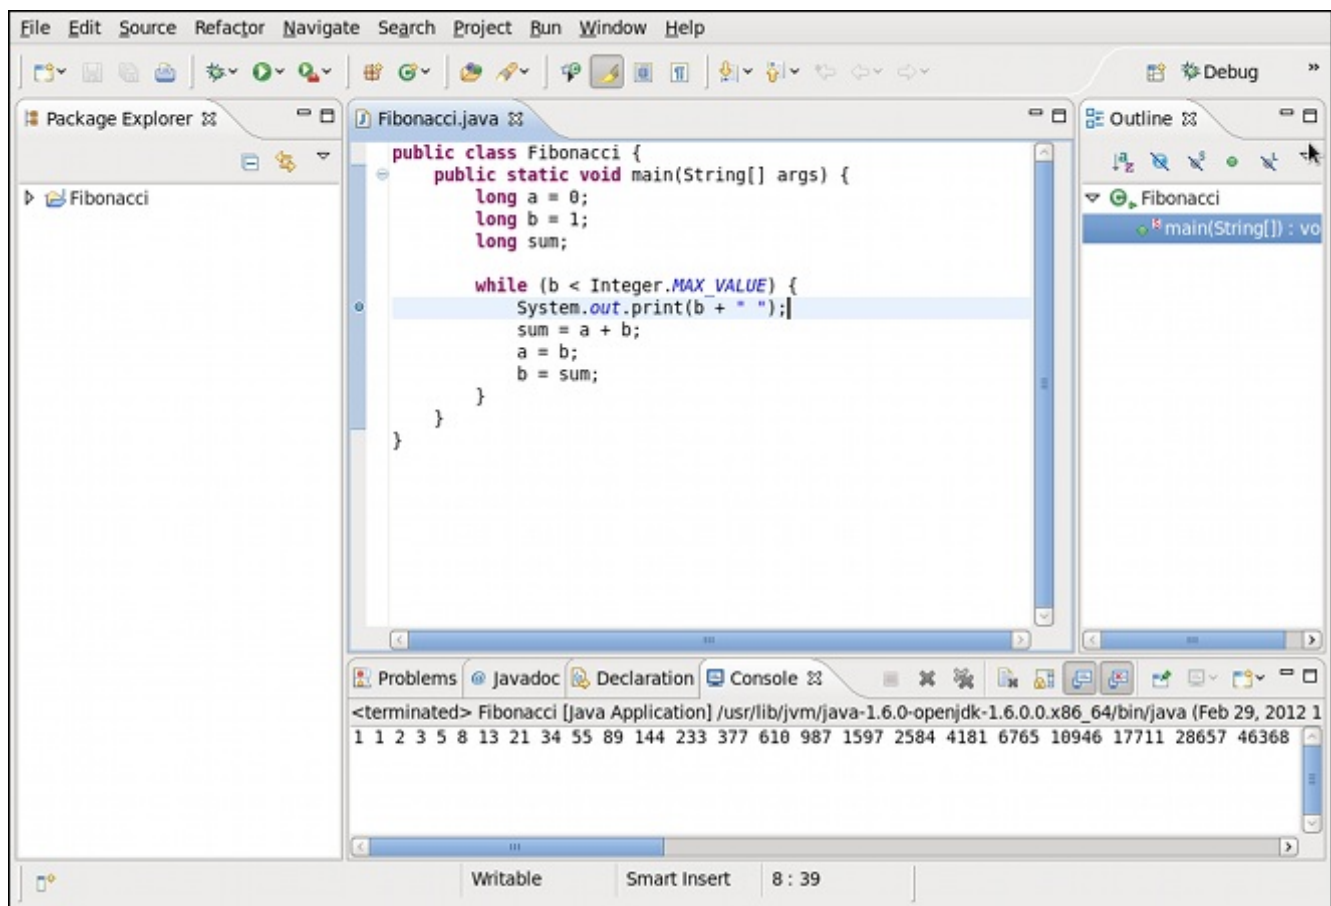


Figure 10.1. Sample Eclipse Session

### 10.1. Installing Eclipse

In Red Hat Enterprise Linux 6, the Eclipse development environment is distributed as a set of several smaller packages such as *eclipse-cdt*, which provides C/C++ support, *eclipse-jdt*, which provides support for Java, or *eclipse-valgrind*, which integrates various Valgrind tools into Eclipse. In addition, a package group named **Eclipse** is available to make it easier to install all packages that form the Eclipse integrated development environment at the same time.

To install the **Eclipse** package group and all its dependencies on your system, type the following at a shell prompt as **root**:

```
yum groupinstall Eclipse
```

This installs *eclipse-cdt*, *eclipse-jdt*, *eclipse-valgrind*, and other packages to the system from the Red Hat Enterprise Linux 6 x86/x86\_64 Supplementary channel. To install all Eclipse packages, run as **root**:

```
yum install eclipse-*
```



### Note

Eclipse is part of Red Hat Enterprise Linux rather than Red Hat Developer Toolset. It is only available for Red Hat Enterprise Linux 6 on 32-bit and 64-bit Intel and AMD architectures, and supports only C and C++ development. The Red Hat Enterprise Linux 6 version does *not* provide support for the Fortran programming language.

## 10.2. Running Eclipse

To start Eclipse with support for the GNU Compiler Collection and GNU Debugger from Red Hat Developer Toolset, type the following at a shell prompt:

```
scl enable devtoolset-1.1 'eclipse'
```

This invokes the Red Hat Enterprise Linux Eclipse IDE using the Developer Toolset **gcc** compiler and **gdb** debugger instead of the Red Hat Enterprise Linux system equivalents.



### Important

If you are working on a project that you previously built with the Red Hat Enterprise Linux version of the GNU Compiler Collection, make sure that you discard all previous build results. To do so, open the project in Eclipse and select **Project** → **Clean** from the menu.

You can execute any command using the **scl** utility, causing it to be run with the Red Hat Developer Toolset binaries used in preference to the Red Hat Enterprise Linux system equivalents. This allows you to run a shell session with Red Hat Developer Toolset **gcc** and **gdb** as the default compiler and debugger for Eclipse:

```
scl enable devtoolset-1.1 'bash'
```



### Note

To verify the version of **gcc** you are using at any point, type the following at a shell prompt:

```
which gcc
```

Red Hat Developer Toolset's **gcc** executable path will begin with **/opt**. Alternatively, you can use the following command to confirm that the version number matches that for Red Hat Developer Toolset **gcc**:

```
gcc -v
```

During its startup, Eclipse prompts you to select a *workspace*, that is, a directory in which you want to



store your projects. Either use `~/workspace/`, which is the default option, or click the **Browse** button to browse your file system and select a custom directory. You can also select the **Use this as the default and do not ask again** check box to prevent Eclipse from prompting you the next time you run it. When you are done, click the **OK** button to confirm the selection and proceed with the startup.

## 10.3. Additional Resources

A detailed description of Eclipse and all its features is beyond the scope of this book. For more information, refer to the resources listed below.

### Installed Documentation

- ▶ Eclipse includes a built-in **Help** system which provides extensive documentation for each integrated feature and tool. This greatly decreases the initial time investment required for new developers to become fluent in its use. The use of this Help section is detailed in the *Red Hat Enterprise Linux Developer Guide* linked below.

### Online Documentation

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides more information on Eclipse, including a description of the user interface, overview of available development toolkits, or instructions on how to use it to build RPM packages.

### See Also

- ▶ [Chapter 1, Red Hat Developer Toolset](#) provides an overview of Red Hat Developer Toolset and more information on how to install it on your system.
- ▶ [Chapter 2, GNU Compiler Collection \(GCC\)](#) provides information on how to compile programs written in C, C++, and Fortran on the command line.

## Part V. Getting Help

## Chapter 11. Accessing Red Hat Product Documentation

**Red Hat Product Documentation** located at <https://access.redhat.com/knowledge/docs/> serves as a central source of information. It is currently translated in 22 languages and for each product, it provides different kinds of books from release and technical notes to installation, user, and reference guides in HTML, PDF, and EPUB formats.

Below is a brief list of documents that are directly or indirectly relevant to this book.

### Red Hat Developer Toolset

- ▶ [Red Hat Developer Toolset 1.1 Release Notes](#) — The *Release Notes* for Red Hat Developer Toolset 1.1 provide more information about this product.
- ▶ [Red Hat Developer Toolset 1.1 Software Collections Guide](#) — The *Software Collections Guide* for Red Hat Developer Toolset 1.1 explains the concept of Software Collections and documents the `sc1` tool.

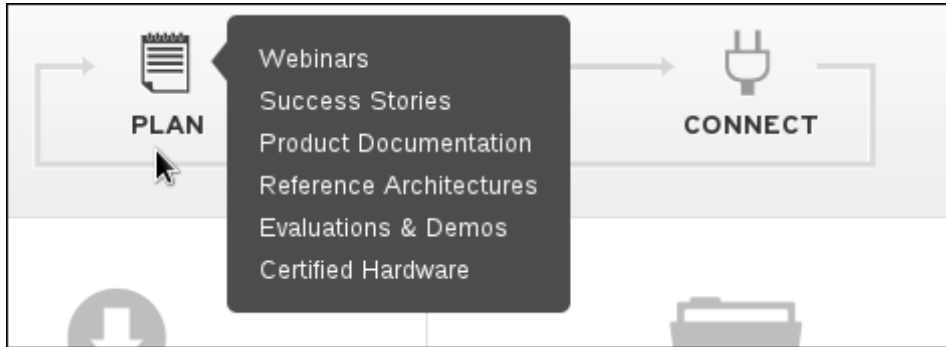
### Red Hat Enterprise Linux

- ▶ [Red Hat Enterprise Linux 6 Developer Guide](#) — The *Developer Guide* for Red Hat Enterprise Linux 6 provides detailed information about libraries and runtime support, compiling and building, debugging, and profiling.
- ▶ [Red Hat Enterprise Linux 6 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 6 explains how to obtain, install, and update the system.
- ▶ [Red Hat Enterprise Linux 5 Installation Guide](#) — The *Installation Guide* for Red Hat Enterprise Linux 5 explains how to obtain, install, and update the system.
- ▶ [Red Hat Enterprise Linux 6 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 6 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 6.
- ▶ [Red Hat Enterprise Linux 5 Deployment Guide](#) — The *Deployment Guide* for Red Hat Enterprise Linux 5 documents relevant information regarding the deployment, configuration, and administration of Red Hat Enterprise Linux 5.

## Chapter 12. Accessing the Customer Portal

The **Customer Portal** is available to all Red Hat subscribers and can be accessed at <https://access.redhat.com/home>. This web page serves as a pointer to a vast number of resources but of most interest to developers are the **Plan**, **Deploy**, and **Connect** menus. These include links to all the resources needed during each stage of the development.

### 12.1. The Plan Menu



**Figure 12.1. The Plan Menu**

In the **Plan** menu are resources needed to plan a development project. These menu items provide proven resources to implement the best solution the first time. This includes:

#### **Webinars**

The Webinars page contains information on upcoming Red Hat and open source events around the world. Here you can register for upcoming webinars or watch archived ones on demand.

#### **Success Stories**

Read the success stories of other Red Hat customers to learn how leading organizations are finding unbeatable value, performance, security and reliability with Red Hat solutions.

#### **Product Documentation**

This provides a list of links to the various Red Hat documents, including books for Red Hat Enterprise Linux, Identity Management and Infrastructure, Red Hat Enterprise Storage, JBoss Enterprise Middleware, and System Management.

#### **Reference Architectures**

Reference Architectures contains a list of whitepapers that detail technical case studies of solutions that have been built, tested, and bench-marked by senior Red Hat engineers. They explain the capabilities and limitations of a given solution, as well as detailed notes on how to implement the solution.

#### **Evaluations & Demos**

You can download free evaluations of various Red Hat products from here, including Red Hat Enterprise Linux, Red Hat Enterprise Virtualization, Red Hat Storage Appliance, and JBoss Enterprise Middleware evaluations.

## Certified Hardware

This section has information on what systems, components, and peripherals Red Hat Enterprise Linux 6, 5, and 4 support.

## 12.2. The Deploy Menu



**Figure 12.2. The Deploy Menu**

In the **Deploy** menu are resources needed to deploy a development project. This includes:

### Getting Started

The Getting Started page has links to information to help get up and running with Red Hat subscriptions, including product registration information, accessing your resources, and engaging in global support. It also has links to a Red Hat Welcome Kit and a Quick Guide to Red Hat Support.

### Downloads

Here is where you can download all that Red Hat offers with descriptions of what each entails.

### Activate a New Subscription

After purchasing a Red Hat subscription this is where you go to activate it. Note that the **Activate a New Subscription** section requires you to enter your Red Hat login and password.

### Support Essentials

Here you can find a list of articles and group discussions, viewable by new posts, most popular, and recent comments, as well as the most recent Red Hat errata.

### Security

Red Hat releases errata to address bugs, provide enhancements, or to fix security vulnerabilities. With each erratum an advisory is supplied to give the details of the issues being fixed, as well as how to obtain and install the required software packages. This section has information about the errata, including:

- ▶ Checking the security update policy and lifetime for all Red Hat products

- ▶ Getting the latest security updates for Red Hat products
- ▶ Getting notified of new security updates
- ▶ Finding out if a specific CVE affects a Red Hat product
- ▶ Reporting a security vulnerability
- ▶ How we measure security vulnerabilities
- ▶ Security Response Team mission and standards of service

## Support

This is where all the information regarding Red Hat support can be found, including links for:

- ▶ Support Cases
- ▶ Support Programs
- ▶ Product Life Cycles
- ▶ Supported Environments
- ▶ Help & Assistance
- ▶ Site Help

## 12.3. The Connect Menu

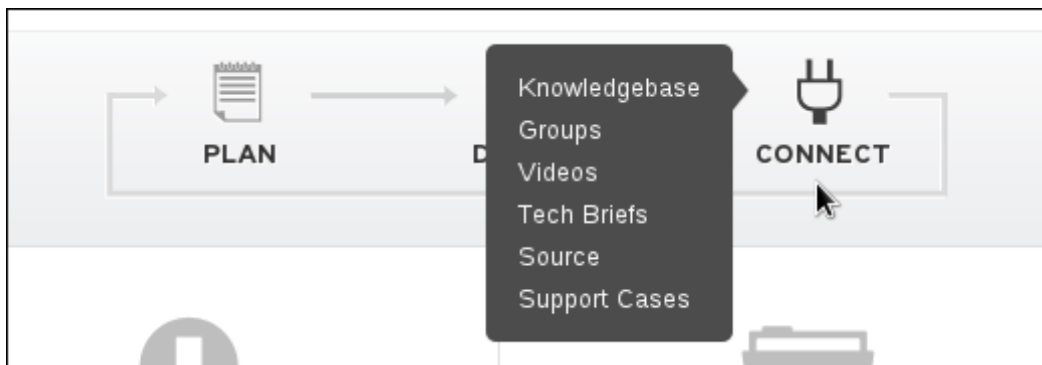


Figure 12.3. The Connect Menu

The **Connect** menu allows you to connect to the industry's best engineers and collaborate with industry peers. This includes:

### Knowledgebase

The Knowledgebase contains a large number of whitepapers on a variety of topics which can serve as another source of documentation for your project. Note that the **Knowledgebase** section requires you to enter your Red Hat login and password.

### Groups

This section allows users to join a group to collaborate on the documents. They can also create and comment on discussions to interact with other customers, Red Hat support staff, and certified engineers. Note that the **Groups** section requires you to enter your Red Hat login and password.

### Videos

A number of videos on how to get vital solutions, useful tips, product demonstrations, and inside information can be accessed from here. You can also rate and comment on all of the videos.

### **Tech Briefs**

Tech briefs provide practical advice to help solve real-world problems with Red Hat products. Each tech brief provides a detailed use case covering best practices, how-to instructions, or detailed discussions on how to use Red Hat technology effectively, and are reviewed and tested by Red Hat engineers.

### **Source**

This is the kernel source browser and contains links to all the kernel sources broken into each individual patch, unlike the kernel srpm which is shipped as one big patch. Note that the **Source** section requires you to enter your Red Hat login and password.

### **Support Cases**

This is where users can view all their support cases, make changes or view any updates. Note that the **Support Cases** section requires you to enter your Red Hat login and password.

## Chapter 13. Contacting Global Support Services

Unless you have a Self-Support subscription, when both the Red Hat Documentation website and Customer Portal fail to provide the answers to your questions, you can contact **Global Support Services (GSS)**.

### 13.1. Gathering Required Information

Several items of information should be gathered before contacting GSS.

#### Background Information

Ensure you have the following background information at hand before calling GSS:

- ▶ Hardware type, make, and model on which the product runs
- ▶ Software version
- ▶ Latest upgrades
- ▶ Any recent changes to the system
- ▶ An explanation of the problem and the symptoms
- ▶ Any messages or significant information about the issue



#### Note

If you ever forget your Red Hat login information, it can be recovered at <https://access.redhat.com/site/help/LoginAssistance.html>.

#### Diagnostics

The diagnostics report for Red Hat Enterprise Linux is required as well. This report is also known as a *sosreport* and the program to create the report is provided by the *sos* package. To install the *sos* package and all its dependencies on your system, type the following at a shell prompt as **root**:

```
yum install sos
```

To generate the report, run as **root**:

```
sosreport
```

For more information, access the Knowledgebase article at <https://access.redhat.com/kb/docs/DOC-3593>.

#### Account and Contact Information

In order to help you, GSS requires your account information to customize their support, as well contact information to get back to you. When you contact GSS ensure you have your:

- ▶ Red Hat customer number or Red Hat Network (RHN) login name
- ▶ Company name
- ▶ Contact name
- ▶ Preferred method of contact (phone or email) and contact information (phone number or email address)



## Issue Severity

Determining an issue's severity is important to allow the GSS team to prioritize their work. There are four levels of severity.

### Severity 1 (urgent)

A problem that severely impacts your use of the software for production purposes. It halts your business operations and has no procedural work around.

### Severity 2 (high)

A problem where the software is functioning but production is severely reduced. It causes a high impact to business operations and no work around exists.

### Severity 3 (medium)

A problem that involves partial, non-critical loss of the use of the software. There is a medium to low impact on your business and business continues to function by utilizing a work around.

### Severity 4 (low)

A general usage question, report of a documentation error, or a recommendation for a future product improvement.

For more information on determining the severity level of an issue, refer to <https://access.redhat.com/support/policy/severity>.

Once the issue severity has been determined, submit a service request through the Customer Portal under the **Connect** option, or at <https://access.redhat.com/support/contact/technicalSupport.html>. Note that you need your Red Hat login details in order to submit service requests.

If the severity is level 1 or 2, then follow up your service request with a phone call. Contact information and business hours are found at <https://access.redhat.com/support/contact/technicalSupport.html>.

If you have a premium subscription, then after hours support is available for Severity 1 and 2 cases.

Turn-around rates for both premium subscriptions and standard subscription can be found at <https://access.redhat.com/support/offerings/production/sla.html>.

## 13.2. Escalating an Issue

If you feel as though an issue is not being handled correctly or adequately, you can escalate it. There are two types of escalation:

### Technical escalation

If an issue is not being resolved appropriately or if you need a more senior resource to attend to it.

### Management escalation

If the issue has become more severe or you believe it requires a higher priority.

More information on escalation, including contacts, is available at [https://access.redhat.com/support/policy/mgt\\_escalation.html](https://access.redhat.com/support/policy/mgt_escalation.html).

### 13.3. Re-opening a Service Request

If more relevant information regarding a closed service request (such as the problem reoccurring), you can re-open it via the Red Hat Customer Portal at [https://access.redhat.com/support/policy/mgt\\_escalation.html](https://access.redhat.com/support/policy/mgt_escalation.html) or by calling your local support center, the details of which can be found at <https://access.redhat.com/support/contact/technicalSupport.html>.



#### Important

In order to re-open a service request, you need the original service request number.

### 13.4. Additional Resources

For more information, refer to the resources listed below.

#### Online Documentation

- ▶ [Getting Started](#) — The *Getting Started* page serves as a starting point for people who purchased a Red Hat subscription, and offers the *Red Hat Welcome Kit* and *Quick Guide to Red Hat Support* for download.
- ▶ [How can a RHEL Self-Support subscription be used?](#) — A Knowledgebase article for customers with a Self-Support subscription.
- ▶ [Red Hat Global Support Services and public mailing lists](#) — A Knowledgebase article that answers frequent questions about public Red Hat mailing lists.

## Changes in Version 1.0

This appendix documents new features and compatibility changes introduced in Red Hat Developer Toolset 1.0.

### A.1. Changes in GCC 4.7.0

Red Hat Developer Toolset 1.0 is distributed with **GCC 4.7.0**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features and compatibility changes in this release. For details on how to use these new features, refer to the [GCC documentation](#).

#### A.1.1. Status and Features

##### A.1.1.1. C++11

GCC 4.7 provides experimental support for building applications compliant with C++11 using the **-std=c++11** or **-std=gnu++11** command line options. However, there is no guarantee for compatibility between C++11 code compiled by different versions of the compiler. Refer to [Section A.1.3.1, "C++ ABI"](#) for details.

The C++ runtime library, **libstdc++**, supports a majority of the C++11 features. However, there is no or only partial support for some features such as certain properties on type traits or regular expressions. For details, refer to the [libstdc++ documentation](#), which also lists implementation-defined behavior.

Support for C++11 **exception\_ptr** and **future** requires changes to the exception handling runtime in the system **libstdc++** package. These changes will be distributed through the normal Z-stream channel. Refer to the *Guidance on Experimental Features in gcc version 4.7.0* section in the *Red Hat Developer Toolset 1.0 Release Notes* for more information.

##### A.1.1.2. C11

GCC 4.7 provides experimental support for some of the features from the C11 revision of the ISO C standard, and in addition to the previous (now deprecated) **-std=c1x** and **-std=gnu1x** command line options, gcc now accepts **-std=c11** and **-std=gnu11**. Note that since this support is experimental, it may change incompatibly in future releases.

Examples for features that are supported are Unicode strings (including the predefined macros **\_\_STDC\_UTF\_16\_\_** and **\_\_STDC\_UTF\_32\_\_**), nonreturning functions (**\_Noreturn** and `<stdnoreturn.h>`), and alignment support (**\_Alignas**, **\_Alignof**, **max\_align\_t**, and `<stdalign.h>`).

##### A.1.1.3. Parallelism and Concurrency

**GCC 4.7** provides improved support for programming parallel applications:

1. The GCC compilers support the OpenMP API specification for parallel programming, version 3.1. Refer to the [OpenMP](#) website for more information on this specification.
2. The C++11 and C11 standards provide programming abstractions for multi-threaded programs. The respective standard libraries include programming abstractions for threads and thread-related features such as locks, condition variables, or futures. These new versions of the standard also define a memory model that precisely specifies the runtime behavior of a multi-threaded program, such as the guarantees provided by compilers and the constraints programmers have to pay attention to when writing multi-threaded programs.

Note that support for the memory model is still experimental (see below for details). For more

information about the status of support for C++11 and C11, refer to [Section A.1.1.1, “C++11”](#) and [Section A.1.1.2, “C11”](#) respectively.

The rest of this section describes two new GCC features in more detail. Both these features make it easier for programmers to handle concurrency (such as when multiple threads do not run truly in parallel but instead have to synchronize concurrent access to shared state), and both provide atomicity for access to memory but differ in their scope, applicability, and complexity of runtime support.

### C++11 Types and GCC Built-ins for Atomic Memory Access

C++11 has support for *atomic types*. Access to memory locations of this type is atomic, and appears as one indivisible access even when other threads access the same memory location concurrently. The atomicity is limited to a single read or write access or one of the other atomic operations supported by such types (for example, two subsequent operations executed on a variable of atomic type are each atomic separately, but do not form one joint atomic operation).

An atomic type is declared as `atomic<T>`, where *T* is the non-atomic base type and must be trivially copyable (for example, `atomic<int>` is an atomic integer). GCC does not yet support any base type *T*, but only those that can be accessed atomically with the atomic instructions offered by the target architecture. This is not a significant limitation in practice, given that atomics are primarily designed to expose hardware primitives in an architecture-independent fashion; pointers and integrals that are not larger than a machine word on the target are supported as base types. Using base types that are not yet supported results in link-time errors.

The code generated for operations on atomic types, including the memory orders, implements the semantics specified in the C++11 standard. However, support for the C++11 memory model is still experimental, and for example GCC might not always preserve data-race freedom when optimizing code.

GCC also supports new built-ins for atomic memory accesses, which follow the design of the memory model and new atomic operations. The former set of synchronization built-ins (that is, those prefixed with `__sync`) are still supported.

### Transactional Memory

*Transactional Memory* (TM) allows programs to declare that a piece of code is supposed to execute as a transaction, that is, virtually atomically and in isolation from other transactions. GCC's transactional memory runtime library, `libitm`, then ensures this atomicity guarantee when executing the compiled program. Compared to atomic memory accesses, it is a higher-level programming abstraction, because it is not limited to single memory locations, does not require special data types for the data it modifies, and because transactions can contain arbitrary code and be nested within other transactions (with some restrictions explained subsequently).

GCC implements transactions as specified in the [Draft Specification for Transactional Language Constructs for C++, version 1.1](#). This draft does not yet specify the language constructs for C, but GCC already supports a C-compatible subset of the constructs when compiling C source code.

The main language constructs are transaction statements and expressions, and are declared by the `__transaction_atomic` or `__transaction_relaxed` keywords followed by a compound statement or expression, respectively. The following example illustrates how to increment a global variable `y` if another variable `x` has a value less than 10:

```
__transaction_atomic { if (x < 10) y++; }
```

This happens atomically even in a multi-threaded execution of the program. In particular, even though the transaction can load `x` and `y` and store to `y`, all these memory accesses are virtually executed as one indivisible step.

Note that in line with the C++11 memory model, programs that use transactions must be free of data races. Transactions are guaranteed to be virtually executed serially in a global total order that is determined by the transactional memory implementation and that is consistent with and contributes to the happens-before order enforced by the rest of the program (that is, transaction semantics are specified based on the C++11 memory model, see the draft specification linked above). Nonetheless, if a program is not data-race-free, then it has undefined behavior. For example, a thread can first initialize some data and then make it publicly accessible by code like this:

```
init(data);  
__transaction_atomic { data_public = true;} // data_public is initially false
```

Another thread can then safely use the data, for instance:

```
__transaction_atomic { if (data_public) use(data); }
```

However, the following code has a data race and thus results in undefined behavior:

```
__transaction_atomic { temp = copy(data); if (data_public) use(temp); }
```

Here, **copy(data)** races with **init(data)** in the initializing thread, because this can be executed even if **data\_public** is not true. Another example for data races is one thread accessing a variable **x** transactionally and another thread accessing it nontransactionally at potentially the same time. Note that the data can be safely reclaimed using code like this (assuming only one thread ever does this):

```
__transaction_atomic { data_public = false; }  
destruct(data);
```

Here, **destruct()** does not race with potential concurrent uses of the data because after the transaction finishes, it is guaranteed that **data\_public** is false and thus data is private. See the specification and the C++11 memory model for more background information about this.

Note that even if transactions are required to virtually execute in a total order, this does not mean that they execute mutually exclusive in time. Transactional memory implementations attempt to run transactions as much in parallel as possible to provide scalable performance.

There are two variants of transactions: *atomic transactions* (**\_\_transaction\_atomic**) and *relaxed transactions* (**\_\_transaction\_relaxed**). The former guarantee atomicity with regard to all other code, but allow only code that is known to not include nontransactional kinds of synchronization, such as atomic or volatile memory access. In contrast, relaxed transactions allow all code (for example calls to I/O functions), but only provide atomicity with regard to other transactions. Therefore, atomic transactions can be nested within other atomic and relaxed transactions, but relaxed transactions can only be nested within other relaxed transactions. Furthermore, relaxed transactions are likely to be executed with less performance, but this depends on the implementation and available hardware.

GCC verifies these restrictions statically at compile time (for example, the requirements on code allowed to be called from within atomic transactions). This has implications for when transactions call functions that are defined within other compilation unit (source file) or within libraries. To enable such cross-compilation-unit calls for transactional code, the respective functions must be marked to contain code that is safe to use from within atomic transactions. Programmers can do so by adding the **transaction\_safe** function attribute to the declarations of these functions and by including this declaration when defining the function. In turn, GCC then verifies that the code in these functions is safe for atomic transactions and generates code accordingly. If the programmer does not follow these constraints and/or steps, compile-time or link-time errors occur. Note that within a compilation unit, GCC detects automatically whether a function is safe for use within transactions, and the attributes therefore typically do not need to be added. See the draft specification linked above for further details.

GCC's transactional memory support is designed in such a way that it does not decrease the performance of programs that do not use transactions, nor the performance of nontransactional code, except due to the normal kinds of interference by concurrent threads that use the same resources such as the CPU.

Transactional Memory support in GCC and **libitm** is still experimental, and both the ABI and API could change in the future if this is required due to the evolution of the specification of the language constructs, or due to implementation requirements. Note that when executing applications built with the **-gnu-tm** command line option, it is currently a prerequisite to also have the appropriate version of the **libitm.so.1** shared library installed. Refer to the *Guidance on Experimental Features in gcc version 4.7.0* section in the *Red Hat Developer Toolset 1.0 Release Notes* for more information.

#### A.1.1.4. Architecture-specific Options

Red Hat Developer Toolset 1.0 is only available for Red Hat Enterprise Linux 5 and 6, both for the 32-bit and 64-bit Intel and AMD architectures. Consequently, the options described below are only relevant to these architectures.

Optimization for several processors is now available through the command line options described in [Table A.1, "Processor Optimization Options"](#).

**Table A.1. Processor Optimization Options**

Option	Description
<b>-march=core2</b> and <b>-mtune=core2</b>	Optimization for Intel Core 2 processors.
<b>-march=corei7</b> and <b>-mtune=corei7</b>	Optimization for Intel Core i3, i5, and i7 processors.
<b>-march=corei7-avx</b> and <b>-mtune=corei7-avx</b>	Optimization for Intel Core i3, i5, and i7 processors with AVX.
<b>-march=core-avx-i</b>	Optimization for the Intel processor code-named Ivy Bridge with RDRND, FSGSBASE, and F16C.
<b>-march=core-avx2</b>	Optimization for a next-generation processor from Intel with AVX2, FMA, BMI, BMI2, and LZCNT.
<b>-march=bdver2</b> and <b>-mtune=bdver2</b>	Optimization for AMD family 15h processors code-named Piledriver.
<b>-march=btver1</b> and <b>-mtune=btver1</b>	Optimization for AMD family 14 processors code-named Bobcat.
<b>-march=bdver1</b> and <b>-mtune=bdver1</b>	Optimization for AMD family 15 processors code-named Bulldozer.

Support for various processor-specific intrinsics and instructions is now available through the command line options described in [Table A.2, "Support for Processor-specific Intrinsics and Instructions"](#).

**Table A.2. Support for Processor-specific Intrinsics and Instructions**

Option	Description
<b>-mavx2</b>	Support for Intel AVX2 intrinsics, built-in functions, and code generation.
<b>-mbmi2</b>	Support for Intel BMI2 intrinsics, built-in functions, and code generation.
<b>-mlzcnt</b>	Implementation and automatic generation of <code>__builtin_clz*</code> using the <code>lzcnt</code> instruction.
<b>-mfma</b>	Support for Intel FMA3 intrinsics and code generation.
<b>-mfsgsbase</b>	Enables the generation of new segment register read/write instructions through dedicated built-ins.
<b>-mrdnd</b>	Support for the Intel <code>rdrnd</code> instruction.
<b>-mf16c</b>	Support for two additional AVX vector conversion instructions.
<b>-mtbm</b>	Support for TBM (Trailing Bit Manipulation) built-in functions and code generation.
<b>-mbmi</b>	Support for AMD's BMI (Bit Manipulation) built-in functions and code generation.
<b>-mcrc32</b>	Support for <code>crc32</code> intrinsics.
<b>-mmovbe</b>	Enables the use of the <code>movbe</code> instruction to implement <code>__builtin_bswap32</code> and <code>__builtin_bswap64</code> .
<b>-mxop, -mfma4, and -mlwp</b>	Support for the XOP, FMA4, and LWP instruction sets for the AMD Orochi processors.
<b>-mabm</b>	Enables the use of the <code>popcnt</code> and <code>lzcnt</code> instructions on AMD processors.
<b>-mpopcnt</b>	Enables the use of the <code>popcnt</code> instruction on both AMD and Intel processors.

When using the x87 floating-point unit, GCC now generates code that conforms to ISO C99 in terms of handling of floating-point excess precision. This can be enabled by **-fexcess-precision=standard** and disabled by **-fexcess-precision=fast**. This feature is enabled by default when using standards conformance options such as **-std=c99**.

#### A.1.1.5. Link-time Optimization

*Link-time optimization* (LTO) is a compilation technique in which GCC generates an internal representation of each compiled input file in addition to the native code, and writes both to the output object file. Subsequently, when several object files are linked together, GCC uses the internal representations of the compiled code to optimize inter-procedurally across all the compilation units. This can potentially improve the performance of the generated code (for example, functions defined in one file can potentially be inlined when called in another file).

To enable LTO, the **-flto** option needs to be specified at both compile time and link time. For further details, including interoperability with linkers and parallel execution of LTO, refer to the documentation for **-flto** in the [GCC 4.7.0 Manual](#). Also note that the internal representation is not a stable interface, so LTO will only apply to code generated by the same version of GCC.



#### Note

Use of Link-time Optimization with debug generation is not yet supported in gcc-4.7 and so use of the **-flto** and the **-g** options together is unsupported in Developer Toolset.

### A.1.1.6. Miscellaneous

**-Ofast** is now supported as a general optimization level. It operates similar to **-O3**, adds options that can yield better-optimized code, but in turn might invalidate standards compliance (for example, **-ffast-math** is enabled by **-Ofast**).

GCC can now inform users about cases in which code generation might be improved by adding attributes such as **const**, **pure**, and **noreturn** to functions declared in header files. Use the **-wsuggest-attribute=[const|pure|noreturn]** command line option to enable this.

Assembler code can now make use of a **goto** feature that allows for jumps to labels in C code.

### A.1.2. Language Compatibility

In this section, we describe the compatibility between the Red Hat Developer Toolset compilers and the Red Hat Enterprise Linux system compilers at the programming-language level (for example, differences in the implementation of language standards such as C99, or changes to the warnings generated by **-Wall**).

Some of the changes are a result of bug fixing, and some old behaviors have been intentionally changed in order to support new standards, or relaxed in standards-conforming ways to facilitate compilation or runtime performance. Some of these changes are not visible to the naked eye and will not cause problems when updating from older versions. However, some of these changes are visible, and can cause grief to users porting to Red Hat Developer Toolset's version of GCC. The following text attempts to identify major issues and suggests solutions.

#### A.1.2.1. C

Constant expressions are now handled by GCC in a way that conforms to C90 and C99. For code expressions that can be transformed into constants by the compiler but are in fact not constant expressions as defined by ISO C, this may cause warnings or errors.

Ill-formed redeclarations of library functions are not accepted anymore by the compiler. In particular, a function with a signature similar to the built-in declaration of a library function (for example, **abort()** or **memcpy()**) must be declared with **extern "C"** to be considered as a redeclaration, otherwise it is ill-formed.

#### Duplicate Member

Consider the following **struct** declaration:

```
struct A { int *a; union { struct { int *a; }; }; };
```

Previously, this declaration used to be diagnosed just by the C++ compiler, now it is diagnosed also by the C compiler. Because of the anonymous unions and structs, there is ambiguity about what **.a** actually refers to and one of the fields therefore needs to be renamed.

#### A.1.2.2. C++

#### Header Dependency Changes

**<iostream>**, **<string>**, and other STL headers that previously included **<unistd.h>** as an implementation detail (to get some feature macros for **gthr\*.h** purposes) no longer do so, because it was a C++ standard violation. This can result in diagnostic output similar to the following:



```
error: 'truncate' was not declared in this scope
error: 'sleep' was not declared in this scope
error: 'pipe' was not declared in this scope
error: there are no arguments to 'offsetof' that depend on a template
parameter, so a declaration of 'offsetof' must be available
```

To fix this, add the following line early in the source or header files that need it:

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

Many of the standard C++ library include files have been edited to no longer include `<cstdlib>` to get namespace-`std`-scoped versions of `size_t` and `ptrdiff_t`. As such, C++ programs that used the macros `NULL` or `offsetof` without including `<cstdlib>` will no longer compile. The diagnostic produced is similar to the following:

```
error: 'ptrdiff_t' does not name a type
error: 'size_t' has not been declared
error: 'NULL' was not declared in this scope
error: there are no arguments to 'offsetof' that depend on a template
parameter, so a declaration of 'offsetof' must be available
```

To fix this issue, add the following line:

```
#include <cstdlib>
```

## Name Lookup Changes

G++ no longer performs an extra unqualified lookup that it incorrectly performed in the past. Instead, it implements the two-phase lookup rules correctly, and an unqualified name used in a template must have an appropriate declaration that:

1. is either in scope at the point of the template's definition, or
2. can be found by argument-dependent lookup at the point of instantiation.

Code that incorrectly depends on a second unqualified lookup at the point of instantiation (such as finding functions declared after the template or in dependent bases) will result in compile-time errors.

In some cases, the diagnostics provided by G++ include hints how to fix the bugs. Consider the following code:

```
template<typename T>
int t(T i)
{
    return f(i);
}

int f(int i)
{
    return i;
}

int main()
{
    return t(1);
}
```

The following diagnostics output will be produced:

```
In instantiation of 'int t(T) [with T = int]'
required from here
error: 'f' was not declared in this scope, and no declarations were found by
argument-dependent lookup at the point of instantiation [-fpermissive]
note: 'int f(int)' declared here, later in the translation unit
```

To correct the error in this example, move the declaration of function `f()` before the definition of template function `t()`. The `-fpermissive` compiler flag turns compile-time errors into warnings and can be used as a temporary workaround.

### Uninitialized const

Consider the following declaration:

```
struct A { int a; A (); };
struct B : public A { };
const B b;
```

An attempt to compile this code now fails with the following error:

```
error: uninitialized const 'b' [-fpermissive]
note: 'const struct B' has no user-provided default constructor
```

This happens, because `B` does not have a user-provided default constructor. Either an initializer needs to be provided, or the default constructor needs to be added.

### Visibility of Template Instantiations

The ELF symbol visibility of a template instantiation is now properly constrained by the visibility of its template arguments. For instance, users that instantiate standard library components like `std::vector` with hidden user defined types such as `struct my_hidden_struct` can now expect hidden visibility for `std::vector<my_hidden_struct>` symbols. As a result, users that compile with the `-fvisibility=hidden` command line option should be aware of the visibility of types included from the library headers used. If the header does not explicitly control symbol visibility, types from those headers will be hidden, along with instantiations that use those types. For instance, consider the following code:

```
#include <vector> // template std::vector has default visibility
#include <ctime> // struct tm has hidden visibility
template class std::vector<tm>; // instantiation has hidden visibility
```

One approach to adjusting the visibility of a library header `<foo.h>` is to create a forwarding header on the `-I` include path consisting of the following:

```
#pragma GCC visibility push(default)
#include_next <foo.h>
#pragma GCC visibility push
```

### User-defined Literal Support

When compiling C++ with the `-std={c++11,c++0x,gnu++11,gnu++0x}` command line option, GCC 4.7.0 unlike older versions supports user-defined literals, which are incompatible with some valid ISO C++03 code. In particular, white space is now needed after a string literal before something that could be a valid user defined literal. Consider the following code:

```
const char *p = "foobar"__TIME__;
```

In C++03, the `__TIME__` macro expands to some string literal and is concatenated with the other one. In C++11, `__TIME__` is not expanded and instead, operator `" __TIME__` is being looked up, which results in a warning like:

```
error: unable to find string literal operator 'operator" __TIME__'
```

This applies to any string literal followed without white space by some macro. To fix this, add some white space between the string literal and the macro name.

### Taking the Address of Temporary

Consider the following code:

```
struct S { S (); int i; };
void bar (S *);
void foo () { bar (&S ()); }
```

Previously, an attempt to compile this code produced a warning message, now it fails with an error. This can be fixed by adding a variable and passing the address of this variable instead of the temporary. The `-fpermissive` compiler flag turns compile-time errors into warnings and can be used as a temporary workaround.

### Miscellaneous

G++ now sets the predefined macro `__cplusplus` to the correct value: `199711L` for C++98/03, and `201103L` for C++11.

G++ now properly re-uses stack space allocated for temporary objects when their lifetime ends, which can significantly lower stack consumption for some C++ functions. As a result of this, some code with undefined behavior will now break.

When an extern declaration within a function does not match a declaration in the enclosing context, G++ now properly declares the name within the namespace of the function rather than the namespace which was open just before the function definition.

G++ now implements the proposed resolution of the C++ standard's core issue 253. Default initialization is allowed if it initializes all subobjects, and code that fails to compile can be fixed by providing an initializer such as:

```
struct A { A(); };
struct B : A { int i; };
const B b = B();
```

Access control is now applied to `typedef` names used in a template, which may cause G++ to reject some ill-formed code that was accepted by earlier releases. The `-fno-access-control` option can be used as a temporary workaround until the code is corrected.

G++ now implements the C++ standard's core issue 176. Previously, G++ did not support using the injected-class-name of a template base class as a type name, and lookup of the name found the declaration of the template in the enclosing scope. Now lookup of the name finds the injected-class-name, which can be used either as a type or as a template, depending on whether or not the name is followed by a template argument list. As a result of this change, some code that was previously accepted may be ill-formed, because:

1. the injected-class-name is not accessible because it is from a private base, or
2. the injected-class-name cannot be used as an argument for a template parameter.

In either of these cases, the code can be fixed by adding a nested-name-specifier to explicitly name the template. The first can be worked around with **-fno-access-control**, the second is only rejected with **-pedantic**.

### A.1.2.3. C/C++ Warnings

GCC 4.7.0 adds a number of new warnings that are either enabled by default, or by using the **-Wall** option. Although these warnings do not result in compilation failure on their own, often **-Wall** is used in conjunction with **-Werror**, causing these warnings to act like errors. This section provides a list of these new or newly enabled warnings. Unless noted otherwise, these warnings apply to both C and C++.

The behavior of the **-Wall** command line option has changed and now includes the new warning flags **-Wunused-but-set-variable** and, with **-Wall -Wextra**, **-Wunused-but-set-parameter**. This may result in new warnings in code that compiled cleanly with previous versions of GCC. For example, consider the following code:

```
void fn (void)
{
    int foo;
    foo = bar (); /* foo is never used. */
}
```

The following diagnostic output will be produced:

```
warning: variable "foo" set but not used [-Wunused-but-set-variable]
```

To fix this issue, first see if the unused variable or parameter can be removed without changing the result or logic of the surrounding code. If not, annotate it with **\_\_attribute\_\_((\_\_unused\_\_))**. As a workaround, you can use the **-Wno-error=unused-but-set-variable** or **-Wno-error=unused-but-set-parameter** command line option.

The **-Wenum-compare** option causes GCC to report a warning when values of different enum types are being compared. Previously, this option only worked for C++ programs, but now it works for C as well. This warning is enabled by **-Wall** and may be avoided by using a type cast.

Casting integers to larger pointer types now causes GCC to display a warning by default. To disable these warnings, use the **-Wno-int-to-pointer-cast** option, which is available for both C and C++.

Conversions between NULL and non-pointer types now cause GCC to report a warning by default. Previously, these warnings were only displayed when explicitly using **-Wconversion**. To disable these warnings, use the new **-Wno-conversion-null** command line option.

GCC can now warn when a class that has virtual functions and a non-virtual destructor is destroyed by using **delete**. This is unsafe to do because the pointer might refer to a base class that does not have a virtual destructor. The warning is enabled by **-Wall** and by a new command line option, **-Wdelete-non-virtual-dtor**.

New **-Wc++11-compat** and **-Wc++0x-compat** options are now available. These options cause GCC to display a warning about C++ constructs whose meaning differs between ISO C++ 1998 and ISO C++ 2011 (such as identifiers in ISO C++ 1998 that are keywords in ISO C++ 2011). This warning is enabled by **-Wall** and enables the **-Wnarrowing** option.

### A.1.3. ABI Compatibility

This section describes compatibility between the Red Hat Developer Toolset compilers and the system

compilers on the *application binary interface* (ABI) level.

### A.1.3.1. C++ ABI

Because the upstream GCC community development does not guarantee C++11 ABI compatibility across major versions of GCC, the same applies to use of C++11 with Red Hat Developer Toolset. Consequently, using the `-std=c++11` option is supported in Red Hat Developer Toolset 1.0 only when all C++ objects compiled with that flag have been built using the same major version of Red Hat Developer Toolset. The mixing of objects, binaries and libraries, built by the Red Hat Enterprise Linux 5 or 6 system toolchain GCC using the `-std=c++0x` or `-std=gnu++0x` flags, with those built with the `-std=c++11` or `-std=gnu++11` flags using the GCC in Red Hat Developer Toolset is explicitly not supported.

As later major versions of Red Hat Developer Toolset may use a later major release of GCC, forward-compatibility of objects, binaries, and libraries built with the `-std=c++11` or `-std=gnu++11` options cannot be guaranteed, and so is not supported.

The default language standard setting for Red Hat Developer Toolset is C++98. Any C++98-compliant binaries or libraries built in this default mode (or explicitly with `-std=c++98`) can be freely mixed with binaries and shared libraries built by the Red Hat Enterprise Linux 5 or 6 system toolchain GCC. Red Hat recommends use of this default `-std=c++98` mode for production software development.



#### Important

Use of C++11 features in your application requires careful consideration of the above ABI compatibility information.

Aside from the C++11 ABI, discussed above, [the Red Hat Enterprise Linux Application Compatibility Specification](#) is unchanged for Developer Toolset. When mixing objects built with Developer Toolset with those built with the Red Hat Enterprise Linux v5.x/v6.x toolchain (particularly `.o/.a` files), the Red Hat Developer Toolset toolchain should be used for any linkage. This ensures any newer library features provided only by Developer Toolset are resolved at link-time.

A new standard mangling for SIMD vector types has been added to avoid name clashes on systems with vectors of varying length. By default the compiler still uses the old mangling, but emits aliases with the new mangling on targets that support strong aliases. `-wabi` will now display a warning about code that uses the old mangling.

### A.1.3.2. Miscellaneous

GCC now optimizes calls to various standard C string functions such as `strlen()`, `strchr()`, `strcpy()`, `strcat()` and `stpcpy()` (as well as their respective `_FORTIFY_SOURCE` variants) by transforming them into custom, faster code. This means that there might be fewer or other calls to those functions than in the original source code. The optimization is enabled by default at `-O2` or higher optimization levels. It is disabled when using `-fno-optimize-strlen` or when optimizing for size.

When compiling for 32-bit GNU/Linux and not optimizing for size, `-fomit-frame-pointer` is now enabled by default. The prior default setting can be chosen by using the `-fno-omit-frame-pointer` command line option.

Floating-point calculations on x86 targets and in strict C99 mode are now compiled by GCC with a stricter standard conformance. This might result in those calculations executing significantly slower. It can be disabled using `-fexcess-precision=fast`.

### A.1.4. Debugging Compatibility

GCC now generates DWARF debugging information that uses more or newer DWARF features than previously. GDB contained in the Developer Toolset can handle these features, but versions of GDB older than 7.0 cannot. GCC can be restricted to only generate debugging information with older DWARF features by using the `-gdwarf-2 -gstrict-dwarf` or `-gdwarf-3 -gstrict-dwarf` options (the latter are handled partially by versions of GDB older than 7.0).

Many tools such as **Valgrind**, **SystemTap**, or third-party debuggers utilize debugging information. It is suggested to use the `-gdwarf-2 -gstrict-dwarf` options with those tools.



#### Note

Use of Link-time Optimization with debug generation is not yet supported in gcc-4.7 and so use of the `-flto` and the `-g` options together is unsupported in Developer Toolset.

### A.1.5. Other Compatibility

GCC is now more strict when parsing command line options, and both **gcc** and **g++** report an error when invalid command line options are used. In particular, when only linking and not compiling code, earlier versions of GCC ignored all options starting with `--`. For example, options accepted by the linker such as `--as-needed` and `--export-dynamic` are not accepted by **gcc** and **g++** anymore, and should now be directed to the linker using `-Wl,--as-needed` or `-Wl,--export-dynamic` if that is intended.

Because of the new link-time optimization feature (see [Section A.1.1.5, “Link-time Optimization”](#)), support for the older intermodule optimization framework has been removed and the `-combine` command line option is not accepted anymore.

## A.2. Changes in GDB 7.4

Red Hat Developer Toolset 1.0 is distributed with **GDB 7.4**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features and compatibility changes in this release. For details on how to use these new features, refer to the [GDB documentation](#).

### A.2.1. New Features

#### Changes Since Red Hat Enterprise Linux 6.2

The following features have been added since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.2:

- ▶ Support for linespecs has been improved (in particular, a more consistent handling of ambiguous linespecs, some support for labels in the program's source, and FILE:LINE support now extends to further linespecs types). Breakpoints are now set on all matching locations in all inferiors and will be updated according to changes in the inferior.
- ▶ New inferior control commands **skip function** and **skip file** have been added. These commands can be used to skip certain functions and files when stepping.
- ▶ The **info threads** command now displays the thread name as set by **prctl** or **pthread\_setname\_np**. In addition, new commands **thread name** and **thread find** have been added. The **thread name** command accepts a name as an argument and can be used to set the name of the current thread. The **thread find** command accepts a regular expression and allows

the user to find threads that match it.

- ▶ GDB now provides support for reading and writing a new **.gdb\_index** section. The command **`gdb-add-index`** can be used to add **.gdb\_index** to a file, which allows GDB to load symbols from that file faster. Note that this feature is already present in Red Hat Enterprise Linux 6.1 and later.
- ▶ The **`watch`** command has been adapted to accept **`-location`** as an optional argument.
- ▶ Two new special values can now be used when specifying the current search path for **`libthread_db`**: **`$sdir`** represents the default system locations of shared libraries, and **`$pdir`** stands for the directory with the **`libthread`** that is used by the application.
- ▶ A new command **`info macros`** has been added. This command accepts **`linespec`** as an optional argument and can be used to display the definitions of macros at that **`linespec`** location. Note that in order to do this, the debugged program must be compiled with the **`-g3`** command line option to have macro information available in it.
- ▶ A new command **`alias`** has been added. This command can be used to create an alias of an existing command.
- ▶ The **`info macro`** command now accepts **`-all`** and **`--`** as valid options.
- ▶ To display a function parameter's entry value (that is, the value at the time of function entry), the suffix **`@entry`** can be added to the parameter. GDB now displays **`@entry`** values in backtraces, if available.
- ▶ New **`set print entry-values`** and **`show print entry-values`** commands have been added. The **`set print entry-values`** command accepts **`both`**, **`compact`**, **`default`**, **`if-needed`**, **`no`**, **`only`**, and **`preferred`** as valid arguments and can be used to enable printing of function arguments at function entry. The **`show print entry-values`** command can be used to determine whether this feature is enabled.
- ▶ New **`set debug entry-values`** and **`show debug entry-values`** commands have been added. The **`set debug entry-values`** command can be used to enable printing of debugging information for determining frame argument values at function entry and virtual tail call frames.
- ▶ **`!command`** has been added as an alias of **`shell command`**.
- ▶ The **`watch`** command now accepts **`mask mask_value`** as an argument. This can be used to create masked watchpoints.
- ▶ New **`set extended-prompt`** and **`show extended-prompt`** commands have been added. The **`set extended-prompt`** command enables support for a defined set of escape sequences that can be used to display various information. The **`show extended-prompt`** command can be used to determine whether the extended prompt is enabled.
- ▶ New **`set basenames-may-differ`** and **`show basenames-may-differ`** commands have been added. The **`set basenames-may-differ`** command enables support for source files with multiple base names. The **`show basenames-may-differ`** command can be used to determine whether this support is enabled. The default option is **`off`** to allow faster GDB operations.
- ▶ A new command line option **`-ix`** (or **`--init-command`**) has been added. This option acts like **`-x`** (or **`--command`**), but is executed before loading the debugged program.
- ▶ A new command line option **`-iex`** (or **`--init-eval-command`**) has been added. This option acts like **`-ex`** (or **`--eval-command`**), but is executed before loading the debugged program.

The following changes have been made to the C++ language support since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.2:

- ▶ When debugging a template instantiation, parameters of the template are now put in scope.

The following changes have been made to the Python scripting support since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.2:

- ▶ The **`register_pretty_printer`** function in module **`gdb.printing`** now takes an optional

**replace** argument.

- ▶ The **maint set python print-stack on|off** command has been deprecated and will be deleted in GDB 7.5. The new command **set python print-stack none|full|message** has replaced it.
- ▶ A prompt substitution hook (**prompt\_hook**) is now available to the Python API.
- ▶ A new Python module **gdb.prompt** has been added to the GDB Python modules library.
- ▶ Python commands and convenience-functions located in **data\_directory/python/gdb/command/** and **data\_directory/python/gdb/function/** are now automatically loaded on GDB start-up.
- ▶ Blocks now provide four new attributes: **global\_block**, **static\_block**, **is\_static**, and **is\_global**.
- ▶ The **gdb.breakpoint** function has been deprecated in favor of **gdb.breakpoints**.
- ▶ A new class **gdb.FinishBreakpoint** is provided.
- ▶ Type objects for **struct** and **union** types now allow access to the fields using standard Python dictionary (mapping) methods.
- ▶ A new event **gdb.new\_objfile** has been added.
- ▶ A new function **deep\_items** has been added to the **gdb.types** module.
- ▶ The function **gdb.Write** now accepts an optional keyword **stream**.
- ▶ Parameters can now be sub-classed in Python, which allows for implementation of the **get\_set\_doc** and **get\_show\_doc** functions.
- ▶ Symbols, Symbol Table, Symbol Table and Line, Object Files, Inferior, Inferior Thread, Blocks, and Block Iterator APIs now have an **is\_valid** method.
- ▶ Breakpoints can now be sub-classed in Python, which allows for implementation of the **stop** function that is executed each time the inferior reaches that breakpoint.
- ▶ A new function **gdb.lookup\_global\_symbol** has been added. This function can be used to look up a global symbol.
- ▶ GDB values in Python are now callable if the value represents a function.
- ▶ A new module **gdb.types** has been added.
- ▶ A new module **gdb.printing** has been added.
- ▶ New commands **info pretty-printers**, **enable pretty-printer**, and **disable pretty-printer** have been added.
- ▶ A new **gdb.parameter("directories")** function call is now available.
- ▶ A new function **gdb.newest\_frame** has been added. This function can be used to return the newest frame in the selected thread.
- ▶ The **gdb.InferiorThread** class now supports a new **name** attribute.
- ▶ Support for inferior events has been added. Python scripts can now add observers in order to be notified of events occurring in the process being debugged.

## Changes Since Red Hat Enterprise Linux 5.8

In addition to the above changes, the following features have been added since the release of GNU Debugger included in Red Hat Enterprise Linux 5.8:

- ▶ For remote targets, debugging of shared libraries is now supported by default.
- ▶ New commands **set observer** and **show observer** have been added. The **set observer** command accepts **on** or **off** as an argument and can be used to allow or disallow the GNU Debugger to affect the execution of the debugged program. Use the **show observer** command to determine whether observer mode is enabled.
- ▶ A new convenience variable **\$\_thread** has been added. This variable stores the number of the



current thread.

- ▶ The **source** command now accepts **-s** as a valid option. This option can be used to search for the script in the source search path regardless of the path in the file name.
- ▶ Support for tracepoints, including fast and static tracepoints, has been added to **gdbserver**.
- ▶ The **--batch** command line option has been adapted to disable pagination and queries.
- ▶ Direct support for the reading and writing byte, word, and double-word x86 general purpose registers such as **\$al** has been added.
- ▶ The **commands** command now accepts a range of breakpoints as an argument.
- ▶ The **rbreak** command now accepts a file name as part of its argument. This can be used to limit the functions selected by the supplied regular expression to those that are defined in the specified file.
- ▶ Support for multi-program (sometimes referred to as multi-executable or multi-exec) debugging has been added. In particular, the GNU Debugger now supports the following commands:
  - The **add-inferior** command can be used to add a new inferior.
  - The **clone-inferior** command can be used to create a copy of an inferior with the same executable loaded.
  - The **remove-inferior** command accepts an inferior ID as an argument and can be used to remove an inferior.
- ▶ Support for trace state variables has been added. In particular, the GNU Debugger now supports the following commands:
  - The **tvariable \$variable\_name [= expression]** command can be used to define or modify a trace state variable.
  - The **info tvariables** command can be used to display a list of currently defined trace state variables and their values.
  - The **delete tvariable \$variable\_name...** command can be used to delete one or more trace state variables.
- ▶ A new **ftrace** has been added. This command accepts a function name, a line number, or an address as an argument, and can be used to define a fast tracepoint at that location.
- ▶ Support for disconnected tracing, trace files, and circular trace buffer has been added.
- ▶ A new **teval** command has been added. This command accepts one or more expressions to evaluate at a tracepoint.
- ▶ The GNU Debugger has been adapted to parse the **0b** prefix of binary numbers exactly the same way as the GNU Compiler Collection.
- ▶ The GNU Debugger now supports the following commands for process record and replay:
  - New commands **set record memory** and **show record memory** have been added. The **set record memory** command accepts **on** or **off** as an argument and can be used to enable or disable stopping the inferior when a memory change of the next instruction cannot be recorded. Use the **show record memory-query** command to determine whether this feature is enabled.
  - A new command **record save** has been added. This command accepts a file name as an argument and can be used to save the execution log to a file.
  - A new command **record restore** has been added. This command accepts a file name as an argument and can be used to restore the execution log from a file.
- ▶ A new command **eval** has been added. This command accepts a format string followed by one or more arguments, transforms it to a command, and then executes it.
- ▶ A new command **save breakpoints** has been added. This command accepts a file name as an argument and can be used to store all currently defined breakpoints to a file. To restore the saved breakpoints from this file, use the **source** command.
- ▶ New commands **may-write-registers**, **set may-write-memory**, **set may-insert-breakpoints**, **set may-insert-tracepoints**, **set may-insert-fast-tracepoints**, and

**set may-interrupt** have been added. All of these commands accept either **on** or **off** as an argument, and can be used to set individual permissions for the target.

- ▶ A new command **main info program-spaces** has been added. This command can be used to display information about currently loaded program spaces.
- ▶ New commands **set remote interrupt-sequence** and **show remote interrupt-sequence** have been added. The **set remote interrupt-sequence** command accepts **Ctrl-C**, **BREAK**, and **BREAK-g** as valid arguments, and can be used to specify which interrupt sequence to send to the remote target in order to interrupt its execution. Use the **show remote interrupt-sequence** to determine the current setting.
- ▶ New commands **set remote interrupt-on-connect** and **show remote interrupt-on-connect** have been added. The **set remote interrupt-on-connect** accepts either **on** or **off** as an argument, and can be used to enable sending an interrupt sequence to the remote target when the GNU Debugger connects to it. Use the **show remote interrupt-on-connect** command to determine whether this feature is enabled.
- ▶ The **set remotebreak** and **show remotebreak** commands have been deprecated and users are advised to use **set remote interrupt-sequence** and **show remote interrupt-sequence** instead.
- ▶ The **disassemble** command has been adapted to accept two arguments in the form of **start,+length**.
- ▶ The **source** command can now be used to read commands from Python scripts.

The following changes have been made to the C++ language support since the release of the GNU Debugger included in Red Hat Enterprise Linux 5.8:

- ▶ Argument-dependent lookup (ADL) now directs function search to the namespaces of its arguments regardless of whether the namespace has been imported.
- ▶ In addition to member operators, the GNU Debugger can now look up operators that are:
  - defined in the global scope,
  - defined in a namespace and imported via the **using** directive,
  - implicitly imported from an anonymous namespace, or
  - the argument-dependent lookup (ADL) operators.
- ▶ Support for printing of static const class members that are initialized in the class definition has been enhanced.
- ▶ Support for importing of namespaces has been added.
- ▶ The C++ expression parser has been adapted to handle the cast operators **static\_cast<>**, **dynamic\_cast<>**, **const\_cast<>**, and **reinterpret\_cast<>**.

The following changes have been made to the Python scripting support since the release of the GNU Debugger in Red Hat Enterprise Linux 5.8:

- ▶ The GNU Debugger is now installed with a new directory located at **/opt/rh/devtoolset-1.1/root/usr/share/gdb/python/**. This directory serves as a standard location for Python scripts written for GDB.
- ▶ The Python API has been adapted to provide access to symbols, symbol tables, program spaces, breakpoints, inferiors, threads, and frame's code blocks. Users are now also allowed to create custom GDB parameters from the API and manipulate them by using the **set** and **show** commands.
- ▶ New functions **gdb.target\_charset**, **gdb.target\_wide\_charset**, **gdb.progspaces**, **gdb.current\_progspace**, and **gdb.string\_to\_argv** have been added.
- ▶ A new exception **gdb.GdbError** has been added.
- ▶ The GNU Debugger now searches pretty-printers in the current program space.
- ▶ The GNU Debugger can now enable or disable pretty-printers individually.

- ▶ The GNU Debugger has been adapted to look for names of Python scripts to automatically load in a special section named `.debug_gdb_scripts`.

## A.2.2. Compatibility Changes

### Changes Since Red Hat Enterprise Linux 6.2

The following changes have been made since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.2:

- ▶ A new command **info auto-load** has been added and can be used to display the status of various automatically loaded files. The **info auto-load gdb-scripts** command lists automatically loaded canned sequences of commands, **info auto-load python-scripts** displays the status of automatically loaded Python scripts, **info auto-load local-gdbinit** displays whether a local `.gdbinit` file in the current working directory is loaded, and **info auto-load libthread-db** displays whether the inferior-specific thread debugging shared library is loaded.
- ▶ New commands **set auto-load** and **show auto-load** have been added and can be used to control automatic loading of files:
  - The **set auto-load gdb-scripts** and **show auto-load gdb-scripts** commands control automatic loading of GDB scripts.
  - The **set auto-load python-scripts** and **show auto-load python-scripts** commands control automatic loading of Python scripts.
  - The **set auto-load local-gdbinit** and **show auto-load local-gdbinit** commands control automatic loading of `.gdbinit` from the current working directory.
  - The **set auto-load libthread-db** and **show auto-load libthread-db** commands control automatic loading of inferior-specific `libthread_db`.
  - The **set auto-load scripts-directory** and **show auto-load scripts-directory** commands control the list of directories from which to automatically load GDB and Python scripts.
  - The **set auto-load safe-path** and **show auto-load safe-path** commands control the list of directories from which it is safe to automatically load all previously mentioned items.
  - The **set debug auto-load** and **show debug auto-load** commands control displaying of debugging information for all previously mentioned items.

The **set auto-load off** command can be used to disable automatic loading globally. You can also use **show auto-load** with no subcommand to display current settings of all previously mentioned items.

- ▶ The **maint set python auto-load on|off** command has been replaced with **set auto-load python-scripts on|off**.
- ▶ The **maintenance print section-scripts** command has been renamed to **info auto-load python-scripts [pattern]** and is no longer classified as a maintenance-only command.
- ▶ Support for the Guile extension language has been removed.
- ▶ The GNU Debugger has been adapted to follow GCC's rules on accessing volatile objects when reading or writing target state during expression evaluation.

### Changes Since Red Hat Enterprise Linux 5.8

No noteworthy changes that affect compatibility have been made in addition to those described above.

## A.3. Changes in binutils 2.22.52

Red Hat Developer Toolset 1.0 is distributed with **binutils 2.22.52**, which provides a number of bug

fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features and compatibility changes in this release. For details on how to use these new features, refer to the [binutils documentation](#).

The GNU assembler (**as**), GNU linker (**ld**), newly added ELF linker (**gold**), and other binary tools that are part of the binutils are now released under the GNU General Public License, version 3.

### A.3.1. GNU Linker

Another ELF linker, **gold**, is now available in addition to **ld**, the existing GNU linker. **gold** is intended to be a drop-in replacement for **ld**, so **ld**'s documentation is intended to be the reference documentation. **gold** supports most of **ld**'s features, except notable ones such as MRI-compatible linker scripts, cross-reference reports (**--cref**), and various other minor options. It also provides significantly improved link time with very large C++ applications.

In Red Hat Developer Toolset, the **gold** linker is not enabled by default. Users can explicitly switch between **ld** and **gold** by using the **alternatives** mechanism.

#### A.3.1.1. New Features

##### Changes Since Red Hat Enterprise Linux 6.2

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.2:

- ▶ A new **INPUT\_SECTION\_FLAGS** keyword has been added to the linker script language. This keyword can be used to select input sections by section header flags.
- ▶ A new **SORT\_BY\_INIT\_PRIORITY** keyword has been added to the linker script language. This keyword can be used to sort sections by numerical value of the GCC *init\_priority* attribute encoded in the section name.

##### Changes Since Red Hat Enterprise Linux 5.8

In addition to the above changes, the following features have been added since the release of the binutils included in Red Hat Enterprise Linux 5.8:

- ▶ GNU/Linux targets now support the **STB\_GNU\_UNIQUE** symbol binding, a GNU extension to the standard set of ELF symbol bindings. The binding is passed on to the dynamic linker, which ensures that in the entire process there is only one symbol with the given name and type in use.



#### Note

The implementation of this feature depends on capabilities only found in newer versions of the **glibc** library. Consequently, this feature is currently available in Red Hat Developer Toolset for Red Hat Enterprise Linux 6.

- ▶ A new command line option **--no-export-dynamic** has been added. This option can be used to undo the effect of the **-E** and **--export-dynamic** options.
- ▶ A new command line option **--warn-alternate-em** has been added. This option can be used to display a warning if an ELF format object file uses an alternate machine code.
- ▶ A new linker script function **REGION\_ALIAS** has been added. This function can be used to create alias names of memory regions.
- ▶ A new command line option **-Ttext-segment address** has been added for ELF-based targets. This option can be used to set the address of the first byte of the text segment.

- ▶ A new linker script command **INSERT** has been added. This command can be used to augment the default script.
- ▶ In a linker script input section, it is now possible to specify a file within an archive by using the **archive:file** syntax.
- ▶ The **--sort-common** command line option now accepts **ascending** and **descending** as optional arguments. This can be used to specify which sorting order to use.
- ▶ A new command line option **--build-id** has been added for ELF-based targets. This option can be used to generate a unique per-binary identifier embedded in a note section.
- ▶ A new command line option **--default-script=file\_name** (or **-dT file\_name**) has been added. This option can be used to specify a replacement for the built-in linker script.
- ▶ A new command line option **-Bsymbolic-functions** has been added. When creating a shared library, this option will cause references to global function symbols to be bound to the definitions with the shared library, if such exist.
- ▶ The new command line options **--dynamic-list-cpp-new** and **--dynamic-list-data** have been added, which can be used to modify the dynamic list.

### A.3.1.2. Compatibility Changes

#### Changes Since Red Hat Enterprise Linux 6.2

The following changes have been made since the release of the binutils included in Red Hat Enterprise Linux 6.2:

- ▶ The **--copy-dt-needed-entries** command line option is no longer enabled by default. Instead, **-no-copy-dt-needed-entries** is now the default option.
- ▶ Evaluation of linker script expressions has been significantly improved. Note that this can negatively affect scripts that rely on undocumented behavior of the old expression evaluation.

#### Changes Since Red Hat Enterprise Linux 5.8

In addition to the above changes, the following changes have been made since the release of the binutils included in Red Hat Enterprise Linux 5.8:

- ▶ The **--add-needed** command line option has been renamed to **--copy-dt-needed-entries** in order to avoid confusion with the **--as-needed** option.
- ▶ For GNU/Linux systems, the linker no longer processes any relocations made against symbols of the **STT\_GNU\_IFUNC** type. Instead, it emits them into the resulting binary for processing by the loader.



#### Note

The implementation of this feature depends on capabilities only found in newer versions of the **glibc** library. Consequently, this feature is currently available in Red Hat Developer Toolset for Red Hat Enterprise Linux 6.

- ▶ The **--as-needed** command line option has been adapted to link in a dynamic library in the following two cases:
  1. if the dynamic library satisfies undefined symbols in regular objects, and
  2. if the dynamic library satisfies undefined symbols in other dynamic libraries unless the library is already found in a **DT\_NEEDED** entry of one of the libraries that are already linked.
- ▶ The **-l:file\_name** command line option now searches the library path for a file name called **file\_name** without adding the **.a** or **.so** file extension.

## A.3.2. GNU Assembler

### A.3.2.1. New Features

#### Changes Since Red Hat Enterprise Linux 6.2

The following features have been added since the release of the binutils included in Red Hat Enterprise Linux 6.2:

- ▶ The GNU Assembler no longer requires double ampersands in macros.
- ▶ A new **--compress-debug-sections** command line option has been added to enable compression of DWARF debug information sections in the relocatable output file. Compressed debug sections are currently supported by the **readelf**, **objdump**, and **gold** tools, but not by **ld**.

#### Changes Since Red Hat Enterprise Linux 5.8

In addition to the above changes, the following features have been added since the release of the binutils included in Red Hat Enterprise Linux 5.8:

- ▶ GNU/Linux targets now support **gnu\_unique\_object** as a value of the **.type** pseudo operation. This value can be used to mark a symbol as globally unique in the entire process.
- ▶ Support for the new discriminator column in the DWARF line table with a discriminator operand for the **.loc** directive has been added.
- ▶ The **.type** pseudo operation now accepts a type of **STT\_GNU\_IFUNC**. This can be used to indicate that if the symbol is the target of a relocation, its value should not be used. Instead, the function should be invoked and its result used as the value.
- ▶ A new pseudo operation **.cfi\_val\_encoded\_addr** has been added. This pseudo operation can be used to record constant addresses in unwind tables without runtime relocation.
- ▶ A new command line option **-mssse-check=[none|error|warning]** has been added for x86 targets.
- ▶ The **-a** command line option now accepts **g** as a valid sub-option. This combination can be used to enable assembly listings with additional information about the assembly, including the list of supplied command line options or the assembler version.
- ▶ A new command line option **-mssse2avx** has been added for x86 targets. This option can be used to encode SSE instructions with VEX prefix.
- ▶ x86 targets now support the Intel XSAVE, EPT, MOVBE, AES, PCLMUL, and AVX/FMA instructions.
- ▶ New command line options **-march=cpu[,+extension...]**, **-mtune=cpu**, **-mmnemonic=[att|intel]**, **-msyntax=[att|intel]**, **-mindex-reg**, **-mnaked-reg**, and **-mold-gcc** have been added for x86 targets.
- ▶ New pseudo operations **.string16**, **.string32**, and **.string64** have been added. These pseudo operations be used to generate wide character strings.
- ▶ The i386 port now supports the SSE5 instruction set.
- ▶ A new pseudo operation **.reloc** has been added. This pseudo operation serves as a low-level interface for creating relocations.

### A.3.2.2. Compatibility Changes

No noteworthy changes that affect compatibility have been made.

## A.3.3. Other Binary Tools

### A.3.3.1. New Features

## Changes Since Red Hat Enterprise Linux 6.2

The following features have been added since the release of the binutils included in Red Hat Enterprise Linux 6.2:

- ▶ The **readelf** and **objdump** tools can now display the contents of the **.debug\_macro** sections.
- ▶ New **--dwarf-start** and **--dwarf-end** command line options have been added to the **readelf** and **objdump** tools. These options are used by the new Emacs mode (see the **dwarf-mode.el** file).
- ▶ A new **--interleave-width** command line option has been added to the **objcopy** tool to allow the use of the **--interleave** to copy a range of bytes from the input to the output.
- ▶ A new **--dyn-syms** command line option has been added to the **readelf** tool. This option can be used to dump dynamic symbol table.
- ▶ A new tool, **elfedit**, has been added to the binutils. This tool can be used to directly manipulate ELF format binaries.
- ▶ A new command line option **--addresses** (or **-a** for short) has been added to the **addr2line** tool. This option can be used to display addresses before function and source file names.
- ▶ A new command line option **--pretty-print** (or **-p** for short) has been added to the **addr2line** tool. This option can be used to produce human-readable output.

## Changes Since Red Hat Enterprise Linux 5.8

In addition to the above changes, the following features have been added since the release of the binutils included in Red Hat Enterprise Linux 5.8:

- ▶ A new command line option **--insn-width=width** has been added to the **objdump** tool. This option can be used to specify the number of bytes to be displayed on a single line when disassembling instructions.
- ▶ A new command line option **--relocated-dump=name|number** has been added to the **readelf** tool. This option can be used to display the relocated contents of a section as a sequence of bytes.
- ▶ A new command line option **--external-symbols-table=filename** has been added to the **gprof** tool. This option can be used to read a symbol table from a certain file.
- ▶ **bfd** now supports a plugin target, which can be used to get basic support for new file formats by having the plugin target load the same shared objects used by **gold**.
- ▶ The **--dwarf** (or **-W** for short) command line option of the **objdump** tool has been adapted to be as flexible as the **--debug-dump** (or **-w**) option of **readelf**.
- ▶ New command line options **--prefix=prefix** and **--prefix-strip=level** have been added to the **objdump** tool. These options can be used to add absolute paths for the **--source** (or **-S** for short) option.
- ▶ A new command line option **-wL** has been added to the **readelf** tool. This option can be used to dump decoded contents of the **.debug\_line** section.
- ▶ “Thin” archives are now supported. Instead of containing object files, such archives contain just pathnames pointing to those files.
- ▶ A new command line option **-F** has been added to the **objdump** tool. This option can be used to include file offsets in the disassembly.
- ▶ A new command line option **-c** has been added to the **readelf** tool. This option can be used to allow string dumps of archive symbol index.
- ▶ The i386 port now supports the SSE5 instruction set.
- ▶ A new command line option **-p** has been added to the **readelf** tool. This option can be used to allow string dumps of sections.

### A.3.3.2. Compatibility Changes

#### Changes Since Red Hat Enterprise Linux 6.2

No noteworthy changes that affect compatibility have been made.

#### Changes Since Red Hat Enterprise Linux 5.8

The following changes have been made since the release of the binutils included in Red Hat Enterprise Linux 5.8:

- ▶ The **--as-needed** command line option has been adapted to link in a dynamic library in the following two cases:
  1. if the dynamic library satisfies undefined symbols in regular objects, and
  2. if the dynamic library satisfies undefined symbols in other dynamic libraries unless the library is already found in a **DT\_NEEDED** entry of one of the libraries that are already linked.



## Changes in Version 1.1

This appendix documents new features and compatibility changes introduced in Red Hat Developer Toolset 1.1. For a complete list of changes that were introduced in the previous version of the product, refer to [Appendix A, Changes in Version 1.0](#).

### B.1. Changes in GCC 4.7.2

Red Hat Developer Toolset 1.1 is distributed with **GCC 4.7.2**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. The main change in this version compared to Red Hat Developer Toolset 1.0 is that Fortran development is now supported. Additionally, support for Fortran 2003 and 2008 has been significantly improved relative to the GCC versions shipped with Red Hat Enterprise Linux (see details below).

#### B.1.1. C++11 Compatibility

An asynchronous erratum advisory, [RHBA-2012-1127](#), that changed the ABI in C++11 mode, was previously published for Red Hat Developer Toolset 1.0. This change is included by default in Red Hat Developer Toolset 1.1. Consequently, any C++11 code compiled with the Red Hat Developer Toolset 1.0 version of GCC *before* applying the erratum might be incompatible with C++11 code compiled by the Red Hat Developer Toolset 1.1 version. There is no change in the C++11 ABI between Red Hat Developer Toolset 1.0 with erratum RHBA-2012-1127 applied and Red Hat Developer Toolset 1.1.

#### B.1.2. Fortran

##### B.1.2.1. New Features

- ▶ A new compile flag **-fstack-arrays** has been added. This flag causes all local arrays to be put on stack memory, which can significantly improve the performance of some programs. Note that programs that use very large local arrays may require you to extend your runtime limits for stack memory.
- ▶ Compile time has been significantly improved. For example, the improvement may be noticeable when working with programs that use large array constructors.
- ▶ To improve code generation and diagnostics, the **-fwhole-file** compile flag is now enabled by default, and can be used with a newly supported **-fwhole-program** flag. To disable it, use the deprecated **-fno-whole-file** flag.
- ▶ A new command line option **-M** is now supported. Similarly to **gcc**, this option allows you to generate Makefile dependencies. Note that the **-cpp** option may be required as well.
- ▶ The **-finit-real=** command line option now supports **snan** as a valid value. This allows you to initialize REAL and COMPLEX variables with a signaling NaN (*not a number*), and requires you to enable trapping (for example, by using the **-ffpe-trap=** command line option). Note that compile-time optimizations may turn a signaling NaN into a quiet NaN.
- ▶ A new command line option **-fcheck=** has been added. This option accepts the following arguments:
  - The **-fcheck=bounds** option is equivalent to the **-fbounds-check** command line option.
  - The **-fcheck=array-temps** option is equivalent to the **-fcheck-array-temporaries** command line option.
  - The **-fcheck=do** option checks for invalid modification of loop iteration variables.
  - The **-fcheck=recursive** option checks for recursive calls to subroutines or functions that are not marked as recursive.
  - The **-fcheck=pointer** option performs pointer association checks in calls, but does not handle undefined pointers nor pointers in expressions.

- The **-fcheck=all** option enables all of the above options.
- ▶ A new command line option **-fno-protect-parens** has been added. This option allows the compiler to reorder REAL and COMPLEX expressions with no regard to parentheses.
- ▶ When OpenMP's **WORKSHARE** is used, array assignments and **WHERE** will now be run in parallel.
- ▶ More Fortran 2003 and Fortran 2008 mathematical functions can now be used as initialization expressions.
- ▶ The **GCC\$** compiler directive now enables support for some extended attributes such as **STDCALL**.

### B.1.2.2. Compatibility Changes

- ▶ The **-Ofast** command line option now automatically enables the **-fno-protect-parens** and **-fstack-arrays** flags.
- ▶ Front-end optimizations can now be disabled by the **-fno-frontend-optimize** option, and selected by the **-ffrontend-optimize** option. The former is essentially only desirable if invalid Fortran source code needs to be compiled (for example, when functions—as compared to subroutines—have side-effects) or to work around compiler bugs.
- ▶ The **GFORTTRAN\_USE\_STDERR** environment variable has been removed, and GNU Fortran now always prints error messages to standard error.
- ▶ The **-fdump-core** command line option and the **GFORTTRAN\_ERROR\_DUMP CORE** environment variable have been removed. When encountering a serious error, GNU Fortran now always aborts the execution of the program.
- ▶ The **-fbacktrace** command line option is now enabled by default. When a fatal error occurs, GNU Fortran now attempts to print a backtrace to standard error before aborting the execution of the program. To disable this behavior, use the **-fno-backtrace** option.
- ▶ GNU Fortran no longer supports the use of the **-M** command line option to generate Makefile dependencies for the module path. To perform this operation, use the **-J** option instead.
- ▶ To significantly reduce the number of warnings, the **-Wconversion** command line option now only displays warnings when a conversion leads to information loss, and a new command line option **-Wconversion-extra** has been added to display warnings about other conversions. The **-Wconversion** option is now enabled with **-Wall**.
- ▶ A new command line option **-Wunused-dummy-argument** has been added. This option can be used to display warnings about unused dummy arguments, and is now enabled with **-Wall**. Note that the **-Wunused-variable** option previously also warned about unused dummy arguments.
- ▶ The **COMMON** default padding has been changed. Previously, the padding was added before a variable. Now it is added after a variable to increase the compatibility with other vendors, as well as to help to obtain the correct output in some cases. Note that this behavior is in contrast with the behavior of the **-falign-commons** option.
- ▶ GNU Fortran no longer links against the **libgfortranbegin** library. The **MAIN\_\_** assembler symbol is the actual Fortran main program and is invoked by the **main** function, which is now generated and put in the same object file as **MAIN\_\_**. Note that the **libgfortranbegin** library is still present for backward compatibility.

### B.1.2.3. Fortran 2003 Features

- ▶ Improved but still experimental support for polymorphism between libraries and programs and for complicated inheritance patterns.
- ▶ Generic interface names which have the same name as derived types are now supported, which allows the creation of constructor functions. Note that Fortran does not support static constructor functions; only default initialization or an explicit structure-constructor initialization are available.
- ▶ Automatic (re)allocation: In intrinsic assignments to allocatable variables, the left-hand side will be automatically allocated (if unallocated) or reallocated (if the shape or type parameter is different). To

avoid the small performance penalty, you can use `a(:) = ...` instead of `a = ...` for arrays and character strings — or disable the feature using `-std=f95` or `-fno-realloc-lhs`.

- ▶ Experimental support of the **ASSOCIATE** construct has been added.
- ▶ In pointer assignments it is now possible to specify the lower bounds of the pointer and, for a rank-1 or a simply contiguous data-target, to remap the bounds.
- ▶ Deferred type parameter: For scalar allocatable and pointer variables the character length can now be deferred.
- ▶ Namelist variables with allocatable attribute, pointer attribute, and with a non-constant length type parameter are now supported.
- ▶ Support has been added for procedure-pointer function results and procedure-pointer components (including **PASS**).
- ▶ Support has been added for allocatable scalars (experimental), **DEFERRED** type-bound procedures, and the **ERRMSG=** argument of the **ALLOCATE** and **DEALLOCATE** statements.
- ▶ The **ALLOCATE** statement now supports type-specs and the **SOURCE=** argument.
- ▶ Rounding (**ROUND=**, **RZ**, ...) for output is now supported.
- ▶ The **INT\_FAST{8,16,32,64,128}\_T** format for **ISO\_C\_BINDING** intrinsic module type parameters is now supported.
- ▶ **OPERATOR(\*)** and **ASSIGNMENT(=)** are now allowed as **GENERIC** type-bound procedures (i.e. as type-bound operators).

#### B.1.2.4. Fortran 2003 Compatibility

Extensible derived types with type-bound procedure or procedure pointer with **PASS** attribute now have to use **CLASS** in line with the Fortran 2003 standard; the workaround to use **TYPE** is no longer supported.

#### B.1.2.5. Fortran 2008 Features

- ▶ A new command line option `-std=f2008ts` has been added. This option enables support for programs that conform to the Fortran 2008 standard and the draft Technical Specification (TS) 29113 on Further Interoperability of Fortran with C. For more information, refer to the [Chart of Fortran TS 29113 Features supported by GNU Fortran](#).
- ▶ The **DO CONCURRENT** construct is now supported. This construct can be used to specify that individual loop iterations do not have any interdependencies.
- ▶ Full single-image support except for polymorphic coarrays has been added, and can be enabled by using the `-fcoarray=single` command line option. Additionally, GNU Fortran now provides preliminary support for multiple images via an MPI-based coarray communication library. Note that the library version is not yet usable as remote coarray access is not yet possible.
- ▶ The **STOP** and **ERROR STOP** statements have been updated to support all constant expressions.
- ▶ The **CONTIGUOUS** attribute is now supported.
- ▶ Use of **ALLOCATE** with the **MOLD** argument is now supported.
- ▶ The **STORAGE\_SIZE** intrinsic inquiry function is now supported.
- ▶ The **NORM2** and **PARITY** intrinsic functions are now supported.
- ▶ The following bit intrinsics have been added:
  - the **POPCNT** and **POPPAR** bit intrinsics for counting the number of 1 bits and returning the parity;
  - the **BGE**, **BGT**, **BLE**, and **BLT** bit intrinsics for bitwise comparisons;
  - the **DSHIFTL** and **DSHIFTR** bit intrinsics for combined left and right shifts;
  - the **MASKL** and **MASKR** bit intrinsics for simple left and right justified masks;
  - the **MERGE\_BITS** bit intrinsic for a bitwise merge using a mask;

- the **SHIFTA**, **SHIFTL**, and **SHIFTR** bit intrinsics for shift operations;
  - the transformational bit intrinsics **IALL**, **IANY**, and **IPARITY**.
- ▶ The **EXECUTE\_COMMAND\_LINE** intrinsic subroutine is now supported.
- ▶ The **IMPURE** attribute for procedures is now supported. This allows the use of **ELEMENTAL** procedures without the restrictions of **PURE**.
- ▶ Null pointers (including **NULL()**) and unallocated variables can now be used as an actual argument to optional non-pointer, non-allocatable dummy arguments, denoting an absent argument.
- ▶ Non-pointer variables with the **TARGET** attribute can now be used as an actual argument to **POINTER** dummies with **INTENT (IN)**.
- ▶ Pointers that include procedure pointers and those in a derived type (pointer components) can now also be initialized by a target instead of only by **NULL**.
- ▶ The **EXIT** statement (with construct-name) can now be used to leave the **ASSOCIATE**, **BLOCK**, **IF**, **SELECT CASE**, and **SELECT TYPE** constructs in addition to **DO**.
- ▶ Internal procedures can now be used as actual arguments.
- ▶ The named constants **INTEGER\_KINDS**, **LOGICAL\_KINDS**, **REAL\_KINDS**, and **CHARACTER\_KINDS** of the intrinsic module **ISO\_FORTRAN\_ENV** have been added. These arrays contain the supported 'kind' values for the respective types.
- ▶ The **C\_SIZEOF** module procedures of the **ISO\_C\_BINDINGS** intrinsic module and the **COMPILER\_VERSION** and **COMPILER\_OPTIONS** module procedures of the **ISO\_FORTRAN\_ENV** intrinsic module have been implemented.
- ▶ The **OPEN** statement now supports the **NEWUNIT=** option. This option returns a unique file unit and therefore prevents inadvertent use of the same unit in different parts of the program.
- ▶ Unlimited format items are now supported.
- ▶ The **INT{8,16,32}** and **REAL{32,64,128}** format for **ISO\_FORTRAN\_ENV** intrinsic module type parameters are now supported.
- ▶ It is now possible to use complex arguments with the **TAN**, **SINH**, **COSH**, **TANH**, **ASIN**, **ACOS**, and **ATAN** functions. Additionally, the new functions **ASINH**, **ACOSH**, and **ATANH** have been added for real and complex arguments, and **ATAN(Y,X)** now serves as an alias for **ATAN2(Y,X)**.
- ▶ The **BLOCK** construct has been implemented.

#### B.1.2.6. Fortran 2008 Compatibility

The implementation of the **ASYNCHRONOUS** attribute in GCC is now compatible with the candidate draft of *TS 29113: Technical Specification on Further Interoperability with C*.

#### B.1.2.7. Fortran 77 Compatibility

When the GNU Fortran compiler is issued with the **-fno-sign-zero** option, the **SIGN** intrinsic now behaves as if zero were always positive.

#### B.1.3. Architecture-specific Options

- ▶ Vectors of type **vector long long** or **vector long** are passed and returned using the same method as other vectors with the **VSX** instruction set. Previously GCC did not adhere to the ABI for 128-bit vectors with 64-bit integer base types (see GCC PR 48857).
- ▶ The **-mrecip** command line option has been added, which indicates whether the reciprocal and reciprocal square root instructions should be used.
- ▶ The **-mvecLibabi=mass** command line option has been added. This can be used to enable the compiler to auto-vectorize mathematical functions using the Mathematical Acceleration Subsystem library.

- ▶ The **-msingle-pic-base** command line option has been added, which instructs the compiler to avoid loading the **PIC** base register in function prologues. The PIC base register must be initialized by the runtime system.
- ▶ The **-mblock-move-inline-limit** command line option has been added, which enables the user to control the maximum size of inlined **memcpy** calls and similar.

## B.2. Changes in GDB 7.5

Red Hat Developer Toolset 1.1 is distributed with **GDB 7.5**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system versions. Below is a comprehensive list of new features in this release. For details on how to use these new features, refer to the [GDB documentation](#).

The following features have been added since the release of the GNU Debugger versions included in Red Hat Enterprise Linux 6.3 and 5.9:

- ▶ The **info os** command has been changed and can now display information on several objects managed by the operating system, in particular:
  - The **info os procggroups** command lists process groups.
  - The **info os files** command lists file descriptors.
  - The **info os sockets** command lists internet-domain sockets.
  - The **info os shm** command lists shared-memory regions.
  - The **info os semaphores** command lists semaphores.
  - The **info os msg** command lists message queues.
  - The **info os modules** command lists loaded kernel modules.
- ▶ GDB now has support for *Static Defined Tracing* (SDT) probes. Currently, the only implemented back end is for SystemTap probes (the **sys/sdt.h** header file). You can set a breakpoint by using the new **-probe**, **-pstap**, or **-probe-stap** options, and inspect the probe arguments by using the new **\$\_probe\_arg** family of convenience variables.
- ▶ The **symbol-reloading** option has been deleted.
- ▶ **gdbserver** now supports STDIO connections, for example:

```
(gdb) target remote | ssh myhost gdbserver - hello
```

- ▶ GDB is now able to print *flag* enums. In a flag enum, all enumerator values have no bits in common when pairwise AND-ed. When GDB prints a value whose type is a flag enum, GDB shows all the constants; for example, for enum **E { ONE = 1, TWO = 2}**:

```
(gdb) print (enum E) 3
$1 = (ONE | TWO)
```

- ▶ The file name part of a linespec now matches trailing components of a source file name. For example, **break gcc/expr.c:1000** now sets a breakpoint in the **build/gcc/expr.c** file, but not in **build/libcpp/expr.c**.
- ▶ The **info proc** and **generate-core-file** commands now work on remote targets connected to **gdbserver**.
- ▶ The command **info catch** has been removed.
- ▶ The Ada-specific **catch exception** and **catch assert** commands now accept conditions at the end of the command.
- ▶ The **info static-tracepoint-marker** command now works on native targets with an in-process agent.

- GDB can now set breakpoints on inline functions.
- The `.gdb_index` section has been updated to include symbols for inline functions. By default, GDB now ignores older `.gdb_index` sections until their `.gdb_index` sections can be recreated. The new command `set use-deprecated-index-sections on` causes GDB to use any older `.gdb_index` sections it finds. If this option is set, the ability to set breakpoints on inline functions is lost in symbol files with older `.gdb_index` sections.
- The `.gdb_index` section has also been updated to record more information about each symbol.
- GDB now provides Ada support for GDB/MI Variable Objects.
- GDB now supports **breakpoint always-inserted mode** in the **record** target.
- **gdbserver** now supports evaluation of breakpoint conditions. Note that you can instruct GDB to send the breakpoint conditions in bytecode form, but **gdbserver** only reports the breakpoint trigger to GDB when its condition evaluates to true.
- The **z0/z1** breakpoint insertion packets have been extended to carry a list of conditional expressions over to the remote stub depending on the condition evaluation mode. You can use the `set remote conditional-breakpoints-packet` command to control the use of this extension.
- A new RSP packet **QProgramSignals** can be used to specify the signals the remote stub can pass to the debugged program without GDB involvement.
- A new command `-info-os` has been added as the MI equivalent of `info os`.
- Output logs, such as `set logging` and related, now include MI output.
- New `set use-deprecated-index-sections on|off` and `show use-deprecated-index-sections on|off` commands have been added. These commands allow you to control the use of deprecated `.gdb_index` sections.
- New `catch load` and `catch unload` commands have been added. These commands allow you to stop execution of a debugged program when a shared library is loaded or unloaded.
- A new command `enable count` has been added. This command allows you to auto-disable a breakpoint after several hits.
- A new command `info vtbl` has been added. This command allows you to show the virtual method tables for C++ and Java objects.
- A new command `explore` has been added. It supports two subcommands `explore value` and `explore type`, and allows you to recursively explore values and types of expressions. Note that this command is only available with Python-enabled GDB.
- A new command `dprintf location,format,args...` has been added. This command allows you to create a dynamic `printf`-type breakpoint, which performs a `printf`-like operation and then resumes program execution.
- New `set print symbol` and `show print symbol` commands have been added. These commands allow you to control whether GDB attempts to display the symbol, if any, that corresponds to addresses it prints. This functionality is enabled by default, but you can restore the previous behavior by running the `set print symbol off` command.
- New `set breakpoint condition-evaluation` and `show breakpoint condition-evaluation` commands have been added. These commands allow you to control whether breakpoint conditions are evaluated by GDB (the `host` option), or by **gdbserver** (the `target` option). The default option, `auto`, chooses the most efficient available mode.
- New `set dprintf-style gdb|call|agent` and `show dprintf-style` commands have been added. These commands allow you to control the way in which a dynamic `printf` is performed: the `gdb` option requests a GDB `printf` command, `call` causes `dprintf` to call a function in the inferior, and `agent` requests that the target agent such as **gdbserver** does the printing.
- New `set dprintf-function expression`, `show dprintf-function`, `set dprintf-channel expression`, and `show dprintf-channel` commands have been added. These

commands allow you to set the function and optional first argument to the call when using the **call** style of dynamic **printf**.

- ▶ New **set disconnected-dprintf on|off** and **show disconnected-dprintf** commands have been added. These commands allow you to control whether agent-style dynamic **printfs** continue to be in effect after GDB disconnects.

The following changes have been made to the Python scripting support since the release of the GNU Debugger included in Red Hat Enterprise Linux 6.3 and 5.9:

- ▶ GDB commands implemented in Python can now be put in the **gdb.COMMAND\_USER** command class.
- ▶ The **maint set python print-stack on|off** command has been removed and replaced by **set python print-stack**.
- ▶ A new class **gdb.printing.FlagEnumerationPrinter** has been added. This class can be used to apply **flag enum**-style pretty-printing to enums.
- ▶ The **gdb.lookup\_symbol** function now works correctly when there is no current frame.
- ▶ The **gdb.Symbol** object now has an additional attribute **line**. This attribute holds the line number in the source at which the symbol was defined.
- ▶ The **gdb.Symbol** object now has an additional attribute **needs\_frame**, and a new method **value**. The **needs\_frame** attribute indicates whether the symbol requires a frame to compute its value, and the **value** method computes the symbol's value.
- ▶ The **gdb.Value** object now has a new method **referenced\_value**. This method can be used to dereference a pointer as well as C++ reference values.
- ▶ The **gdb.Symtab** object now has two new methods, **global\_block** and **static\_block**. These methods return the global and static blocks (as **gdb.Block** objects) of the underlying symbol table respectively.
- ▶ A new method **gdb.find\_pc\_line** returns the **gdb.Symtab\_and\_line** object associated with a PC value.
- ▶ The **gdb.Symtab\_and\_line** object now has an additional attribute **last**. This attribute holds the end of the address range occupied by the code for the current source line.

## B.3. Changes in binutils 2.23.51

Red Hat Developer Toolset 1.1 is distributed with **binutils 2.23.51**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features in this release.

### B.3.1. GNU Linker

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.3 and 5.9:

- ▶ A new **SORT\_NONE** keyword has been added to the linker script language. This keyword can be used to disable section sorting.
- ▶ A new linker-provided symbol, **\_\_ehdr\_start**, has been added. When producing ELF output, this symbol points to the ELF file header (and nearby program headers) in the program's memory image.

### B.3.2. GNU Assembler

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.3 and 5.9:

- ▶ Support for **.bundle\_align\_mode**, **.bundle\_lock**, and **.bundle\_unlock** directives for x86

targets has been added..

- ▶ On x86 architectures, the GNU Assembler now allows **rep bsf**, **rep bsr**, and **rep ret** syntax.

### B.3.3. Other Binary Tools

The following features have been added since the release of binutils included in Red Hat Enterprise Linux 6.3 and 5.9:

- ▶ Support for **dwz -m** optimized debug information has been added.
- ▶ The *devtoolset-1.1-binutils-devel* package now provides the **demangle.h** header file.

## B.4. Changes in elfutils 0.154

Red Hat Developer Toolset 1.1 is distributed with **elfutils 0.154**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features in this release.

### Changes Since Red Hat Enterprise Linux 6.3

The following features have been added since the release of elfutils included in Red Hat Enterprise Linux 6.3:

- ▶ **libdw** now handles compressed debuginfo sections. The **dwarf\_highpc** function now handles the DWARF 4 **DW\_AT\_high\_pc** constant form.
- ▶ **eu-elflint** now accepts executables produced by the **gold** linker (see [Section A.3, “Changes in binutils 2.22.52”](#)).
- ▶ **eu-nm** now supports C++ demangling.
- ▶ **eu-ar** now supports a new modifier **D** for *deterministic output* with no UID, GID, or mtime info. The **U** modifier is the inverse.
- ▶ **eu-readelf** can now print SDT ELF notes (*SystemTap probes*) and the **.gdb\_index** GDB section. It can also now print **DW\_OP\_GNU\_entry\_value** and **DW\_AT\_GNU\_call\_site** families of DIE attributes.
- ▶ **eu-strip** now recognizes a new command line option, **--reloc-debug-sections**.

### Changes Since Red Hat Enterprise Linux 5.9

In addition to the above changes, the following features have been added since the release of elfutils included in Red Hat Enterprise Linux 5.9:

- ▶ DWARF 4 is now supported. Support for **DW\_OP\_GNU\_implicit\_pointer**, **STB\_GNU\_UNIQUE** was added.
- ▶ **libdwfl** now supports automatic decompression of files compressed with **gzip**, **bzip2** and **lzma**, and of Linux kernel images made with **gzip**, **bzip2** or **lzma**. Files named with compression suffixes are searched for Linux kernel images. Core file support was improved. Support has been added for decoding DWARF CFI into location description form. Support has been added for some new DWARF 3 expression operations, which were previously omitted. A new function, **dwfl\_dwarf\_line**, has been added.
- ▶ **eu-elfcmp** has a new flag **--ignore-build-id** to ignore differing build ID bits. The new flag **-l/--verbose** prints all differences.
- ▶ **eu-strip** now recognizes a new flag **--strip-sections** to remove section headers entirely.
- ▶ **libdw** has new functions **dwarf\_next\_unit**, **dwarf\_offdie\_types**, **dwarf\_lineisa**, **dwarf\_linediscriminator**, **dwarf\_lineop\_index**,



**dwarf\_getlocation\_implicit\_pointer**, and **dwarf\_aggregate\_size**.

- ▶ **eu-addr2line** recognizes **-F/--flags** to print more DWARF line information details.
- ▶ **libelf** now supports using more than 65536 program headers in a file. A new function **elf\_getphdrnum** has been added.
- ▶ **eu-addr2line** now accepts the **-j/--section=name** option in the interests of binutils compatibility.
- ▶ **libcpu** Intel SSE4 disassembler support has been added.
- ▶ **eu-readelf** now implements call frame information and exception handling dumping. The **-e** command line option has been added; this is enabled implicitly by **-a**.

## B.5. Changes in Valgrind 3.8.1

Red Hat Developer Toolset 1.1 is distributed with **Valgrind 3.8.1**, which provides a number of bug fixes and feature enhancements over the Red Hat Enterprise Linux system version. Below is a comprehensive list of new features in this release.

### Changes Since Red Hat Enterprise Linux 6.3

The following features have been added since the release of Valgrind included in Red Hat Enterprise Linux 6.3:

- ▶ Support for Intel *Advanced Vector Extensions* (AVX) and **Advanced Encryption Standard** (AES) instruction sets has been added. Note that this support is only available for 64-bit code.
- ▶ A new command line option **--soname-synonyms** has been added. This option can be used to inform **Memcheck**, **Massif**, **Helgrind**, and **DRD** (that is, the tools that replace the **malloc** function) that the implementation of **malloc** is either linked statically into the executable, or is provided by a shared library other than **libc.so**. As a result, users can now process both statically linked programs and programs that use other **malloc** libraries, such as TCMalloc or JEMalloc. This option is, however, very general and does not necessarily apply only to **malloc**-family functions.
- ▶ A new command line option **--redzone-size=number** has been added to tools that provide their own replacement for **malloc** and other functions. This option allows you to specify the size of the padding blocks (also known as *redzones*) that are added before and after each client allocated block. Smaller values reduce the amount of memory required by Valgrind; bigger numbers increase the chance to detect blocks that are overrun or underrun.
- ▶ The **leak\_check** GDB server monitor command now can now be used to control the maximum number of loss records that **Memcheck** outputs.
- ▶ **Memcheck** now uses less memory when working with applications that allocate a large number of blocks or that have many partially-defined bytes.
- ▶ A new GDB server monitor command **block\_list** has been added. This allows **Memcheck** to list the addresses and sizes of the blocks of a leak search loss record.
- ▶ A new GDB server monitor command **who\_points\_at** has been added. This allows **Memcheck** to list the locations that are pointing at a block.
- ▶ When the **Memcheck** tool is executed with a redzone size greater than 0, **VALGRIND\_MALLOCLIKE\_BLOCK** detects an invalid access of these redzones by marking them **noaccess**. Similarly, if a redzone size is given for a memory pool, **VALGRIND\_MEMPOOL\_ALLOC** marks the redzones **noaccess**. Note that this functionality still allows you to find some bugs if the you have forgotten to mark the pool superblock **noaccess**.
- ▶ The performance of **Memcheck**'s memory leak check has been improved. The improvement is especially significant in cases where there are many leaked blocks or many suppression rules used to suppress leak reports.
- ▶ More facilities that can help finding the cause of a data race have been added to **DRD**, namely the --

**ptrace-addr** command line option and the **DRD\_STOP\_TRACING\_VAR(x)** macro.

- ▶ The C++ demangler has been updated to support C++ programs compiled by **g++** up to at least version 4.6.
- ▶ A new command line option **--fair-sched=yes** has been added. This option allows you to select round-robin scheduling of threads, improves responsiveness of interactive multithreaded applications, and improves repeatability of results from the thread checkers Helgrind and DRD. The pipe-based thread serialization mechanism is selected by default.
- ▶ Support for running Valgrind on Valgrind has been improved. Tool developers can now use Valgrind on **Helgrind** or **Memcheck**.
- ▶ **gdbserver** has been adapted to display the float shadow registers as integer rather than float values, as the shadow values are mostly used as bit patterns.
- ▶ The maximum limit for the **--num-callers** command line option has been increased to **500**.
- ▶ The performance of error matching when there are many suppression records in use has been improved.
- ▶ Support for DWARF4 debugging information has been improved.
- ▶ Support for DWARF debug information compressed with **dwz** has been added.
- ▶ Control over the IR optimizer's handling of the tradeoff between performance and precision of exceptions has been improved. In particular, the **--vex-iropt-precise-memory-exns** command line option has been replaced by **--vex-iropt-register-updates**, which provides extended functionality and allows the Valgrind **gdbserver** to always show up-to-date register values to GDB.
- ▶ Modest performance gains have been achieved by the use of translation chaining for JIT-generated code.
- ▶ The memory consumption of the **Memcheck** tool has been reduced in some scenarios.
- ▶ Handling of freed memory in the **Memcheck** tool has been improved, which allows the tool to detect use-after-free that would previously have been missed.
- ▶ The performance of the **Helgrind** tool has been improved. As well, the tool now consumes less memory, especially when working with large, long running applications that perform many synchronization events such as lock, unlock, and similar.
- ▶ **Helgrind** now displays locksets for both threads involved in a race.
- ▶ **Helgrind** now displays error messages that are better formatted and generally easier to understand.
- ▶ **Helgrind** now provides facilities and documentation regarding annotation of thread-safe reference-counted C++ classes.
- ▶ A new command line option **--check-stack-refs=no|yes** has been added to **Helgrind**. This option allows you to disable race checking on thread stacks. The option is set to **yes** by default.
- ▶ A new command line option **--free-is-write=no|yes** has been added to **Helgrind**. This option allows you to enable detection of races where one thread accesses heap memory, but another one frees it without any coordinating synchronization event. The option is set to **no** by default.
- ▶ **DRD** can now produce XML output.
- ▶ Support for delayed thread deletion has been added to **DRD** to detect race conditions that occur close to the end of a thread (the **--join-list-vol** option).
- ▶ The **exp-ptrcheck** tool has been renamed to **exp-sgcheck** (Stack and global Array Checking).
- ▶ The **exp-sgcheck** tool has been reduced in scope to improve performance and remove checking that **Memcheck** does better. In particular, the ability to check for overruns for stack and global arrays is unchanged, but the ability to check for overruns of heap blocks has been removed.
- ▶ Valgrind now provides an embedded GDB server, which allows you to control a Valgrind run from GDB and perform GDB operations such as single stepping, breakpoints, examining data, and so on. Tool-specific functionality is also available, and it is now possible to query the definedness state of variables or memory from within GDB when running **Memcheck**, add arbitrarily large memory watchpoints, etc. To use the GDB server, start Valgrind with the **--vgdb-error=0** command line option and follow the on-screen instructions.

- ▶ Support for unfriendly self-modifying code has been improved. A new command line option **--smc-check=all-non-file** has been added to add the relevant consistency checks only to code that originates in non-file-backed mappings. In effect, this confines the consistency checking only to code that is or might be JIT-generated, and avoids checks on code that must have been compiled ahead of time. Consequently, this functionality significantly improves the performance of applications that generate code at run time.
- ▶ New client requests **VALGRIND\_DISABLE\_ERROR\_REPORTING** and **VALGRIND\_ENABLE\_ERROR\_REPORTING** have been added. These requests enable and disable error reporting on a per-thread and nestable basis, which is useful for hiding errors in particularly troublesome pieces of code. The MPI wrapper library, **libmpiwrap.c**, now uses this facility.
- ▶ The **--mod-funcname** command line option has been added to **cg\_diff**.

## Changes Since Red Hat Enterprise Linux 5.9

In addition to the above changes, the following features have been added since the release of Valgrind included in Red Hat Enterprise Linux 5.9:

- ▶ When running in 64-bit mode, Valgrind now supports the SSE4.2 instruction set with the exception of SSE4.2 AES instructions. In 32-bit mode, Valgrind only provides support up to and including the SSSE3 instruction set.
- ▶ A new processing script **cg\_diff** has been added to **Cachegrind**. This processing script can find the difference between two profiles, and can therefore be used to evaluate the performance effects of a change in a program.
- ▶ The behavior of the **cg\_annotate**'s rarely-used **--threshold** option has been changed.
- ▶ **Callgrind** now supports branch prediction simulation and can optionally count the number of executed global bus events. Note that in order to use this functionality for a better approximation of a "Cycle Estimation" as a derived event, you must manually update the event formula in **KCachegrind**.
- ▶ To accommodate machines with three levels of caches, both **Cachegrind** and **Callgrind** now refer to the LL (last-level) cache rather than the L2 cache. When **Cachegrind** or **Callgrind** auto-detects the cache configuration of such a machine, it now runs the simulation as if the L2 cache is not present. Consequently, the results are less likely to match the true result for the machine and should not be considered authoritative, but provide a general idea about a program's locality.
- ▶ A new command line option **--pages-as-heap** has been added to **Massif**. When this option is enabled, **Massif** tracks memory allocations at the level of memory pages (as mapped by **mmap**, **brk** and similar functions) instead of tracking allocations at the level of heap blocks (as allocated with **malloc**, **new**, or **new[ ]**). Each mapped page is treated as its own block. Interpreting the page-level output is harder than the heap-level output, but allows you to account for every byte of memory used by a program. By default, the **--pages-as-heap** option is disabled.
- ▶ New command line options **--free-is-write** and **--trace-alloc** have been added to **DRD**. The **--free-is-write** option allows you to detect reading from already freed memory, the **--trace-alloc** can be used to trace of all memory allocations and deallocations.
- ▶ A number of new allocations have been added to **DRD**. As well, you can now annotate custom barrier implementations and benign races on static variables.
- ▶ **DRD**'s happens before and happens after annotations have been enhanced and can be used to annotate, for example, a smart pointer implementation.
- ▶ **Helgrind**'s annotation set has been significantly improved to provide a general set of annotations to describe locks, semaphores, barriers, and condition variables. In addition, **Helgrind** now supports annotations to describe thread-safe reference counted heap objects.
- ▶ A new command line option **--show-possibly-lost** has been added to **Memcheck**. By default this option is enabled and causes the leak detector to show possibly-lost blocks.
- ▶ A new experimental heap profiler, **DHAT** (Dynamic Heap Analysis Tool), has been added. **DHAT** keeps track of allocated heap blocks, and also inspects every memory reference to see which block

(if any) is being accessed. This gives a lot of insight into block lifetimes, utilization, turnover, liveness, and the location of hot and cold fields. You can use **DHAT** to do hot-field profiling.

- ▶ Support for unfriendly self-modifying code has been improved, and the extra overhead incurred by **--smc-check=all** has been reduced by approximately a factor of 5 as compared with the previous version of Valgrind.
- ▶ A new command line option **--fullpath-after** has been added. This option can be used to display directory names for source files in error messages, and is combined with a flexible mechanism for specifying which parts of the paths should be shown.
- ▶ A new command line option **--require-text-symbol** has been added. This option stops the execution if a specified symbol is not found in a given shared object when loaded into the process. As a result, working with function intercepting and wrapping is now safer and more reliable.
- ▶ Valgrind now implements more reliable stack unwinding on amd64-linux, particularly in the presence of function wrappers, and with gcc-4.5 compiled code.
- ▶ Valgrind now implements modest scalability (performance improvements) for very large, long-running applications.
- ▶ Valgrind now provides improved support for analyzing programs that are running in **Wine**. Users can now include the **valgrind/valgrind.h**, **valgrind/memcheck.h**, and **valgrind/drd.h** header files in Windows programs that are compiled with MinGW or one of the Microsoft Visual Studio compilers.

## B.6. Changes in OProfile 0.154

Red Hat Developer Toolset 1.1 is distributed with **OProfile 0.9.7**. Below is a comprehensive list of new features in this release.

### Changes Since Red Hat Enterprise Linux 6.3

No noteworthy changes have been made.

### Changes Since Red Hat Enterprise Linux 5.9

The following features have been added since the release of OProfile included in Red Hat Enterprise Linux 5.9:

- ▶ Support has been added for AMD family11h, family12h, family14h, and AMD family15h processors.
- ▶ Generation of XML output has been corrected.
- ▶ Handling of the **--session-dir** command line option has been improved and a possible buffer overflow in the XML generator has been fixed.

## Revision History

<b>Revision 1.1-0</b>	<b>Tue Jan 22 2013</b>	<b>Jaromír Hradílek</b>
Red Hat Developer Toolset 1.1 release of the <i>User Guide</i> .		
<b>Revision 1.0-2</b>	<b>Fri Nov 09 2012</b>	<b>Jaromír Hradílek</b>
Red Hat Developer Toolset 1.1 Beta-2 release of the <i>User Guide</i> .		
<b>Revision 1.0-1</b>	<b>Wed Oct 10 2012</b>	<b>Jaromír Hradílek</b>
Red Hat Developer Toolset 1.1 Beta-1 release of the <i>User Guide</i> .		
<b>Revision 1.0-0</b>	<b>Tue Jun 26 2012</b>	<b>Jaromír Hradílek</b>
Red Hat Developer Toolset 1.0 release of the <i>User Guide</i> .		
<b>Revision 0.0-2</b>	<b>Tue May 22 2012</b>	<b>Jaromír Hradílek</b>
Red Hat Developer Toolset 1.0 Beta-2 release of the <i>User Guide</i> .		
<b>Revision 0.0-1</b>	<b>Tue May 1 2012</b>	<b>Jaromír Hradílek</b>
Red Hat Developer Toolset 1.0 Beta-1 release of the <i>User Guide</i> .		
<b>Revision 0.0-0</b>	<b>Mon Apr 23 2012</b>	<b>Jaromír Hradílek</b>
Initial creation of the book by Publican.		

## Index

### A

#### ABI

- compatibility, [ABI Compatibility](#)

#### addr2line

- features, [New Features](#)  
- overview, [binutils](#)  
- usage, [Using Other Binary Tools](#)

#### application binary interface (see ABI)

#### ar

- overview, [binutils](#)  
- usage, [Using Other Binary Tools](#)

#### as (see GNU assembler)

#### assembling (see GNU assembler)

### B

#### bfd

- features, [New Features](#)

#### binutils, [binutils](#)

- compatibility, [Changes in binutils 2.22.52](#)
- documentation, [Additional Resources](#)
- features, [Main Features](#), [Changes in binutils 2.22.52](#)
- installation, [Installing binutils](#)
- overview, [binutils](#)
- usage, [Using the GNU Assembler](#), [Using the GNU Linker](#), [Using Other Binary Tools](#)
- version, [Main Features](#), [binutils](#), [Changes in binutils 2.22.52](#), [Changes in binutils 2.23.51](#)

## C

### **C programming language**

- compiling, [Using the C Compiler](#), [Preparing a Program for Debugging](#)
- running, [Running a C Program](#)
- support, [GNU C Compiler](#)

### **C++ programming language**

- compiling, [Using the C++ Compiler](#), [Preparing a Program for Debugging](#)
- running, [Running a C++ Program](#)
- support, [GNU C++ Compiler](#)

### **C++11 (see GNU Compiler Collection)**

#### **c++filt**

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

### **C11 (see GNU Compiler Collection)**

#### **Cachegrind**

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

#### **Callgrind**

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

#### **compatibility**

- GNU Compiler Collection, [Language Compatibility](#), [ABI Compatibility](#), [Debugging Compatibility](#), [Other Compatibility](#), [C++11 Compatibility](#), [Compatibility Changes](#), [Fortran 2003 Compatibility](#), [Fortran 2008 Compatibility](#), [Fortran 77 Compatibility](#)
- GNU Debugger, [Compatibility Changes](#)
- Red Hat Developer Toolset, [Compatibility](#)

#### **compiling (see GNU Compiler Collection)**

#### **Customer Portal, [Accessing the Customer Portal](#)**

- Connect menu, [The Connect Menu](#)
- Deploy menu, [The Deploy Menu](#)
- overview, [Accessing the Customer Portal](#)
- Plain menu, [The Plan Menu](#)

## D

**debugging (see GNU Debugger)**

**Developer Toolset (see Red Hat Developer Toolset)**

**documentation**

- Red Hat Product Documentation, [Accessing Red Hat Product Documentation](#)

**DRD**

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

**dwz, [dwz](#)**

- documentation, [Additional Resources](#)
- installation, [Installing dwz](#)
- overview, [dwz](#)
- usage, [Using dwz](#)
- version, [Main Features](#), [dwz](#)

## E

**Eclipse, [Eclipse](#)**

- configuration, [Running Eclipse](#)
- documentation, [Additional Resources](#)
- installation, [Installing Eclipse](#)
- overview, [Eclipse](#)
- usage, [Running Eclipse](#)

**elfedit**

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

**elfutils, [elfutils](#)**

- documentation, [Additional Resources](#)
- installation, [Installing elfutils](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)
- version, [Main Features](#), [elfutils](#), [Changes in elfutils 0.154](#)

**eu-addr2line**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-ar**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)

- usage, [Using elfutils](#)

**eu-elfcmp**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-elflint**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-findtextrel**

- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-make-debug-archive**

- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-nm**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-objdump**

- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-ranlib**

- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-readelf**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-size**

- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-strings**

- overview, [elfutils](#)
- usage, [Using elfutils](#)



**eu-strip**

- features, [Changes in elfutils 0.154](#)
- overview, [elfutils](#)
- usage, [Using elfutils](#)

**eu-unstrip**

- overview, [elfutils](#)
- usage, [Using elfutils](#)

**F****feedback**

- contact information for this manual, [We Need Feedback!](#)

**Fortran programming language**

- compiling, [Using the Fortran Compiler](#)
- running, [Running a Fortran Program](#)
- support, [GNU Fortran Compiler](#)

**G****g++ (see GNU Compiler Collection)****GAS (see GNU assembler)****GCC (see GNU Compiler Collection)****gcc (see GNU Compiler Collection)****GDB (see GNU Debugger)****gfortran (see GNU Compiler Collection)****Global Support Services**

- contacting, [Contacting Global Support Services](#)

**GNU assembler, [Using the GNU Assembler](#)**

- compatibility, [Compatibility Changes](#)
- documentation, [Additional Resources](#)
- features, [New Features, GNU Assembler](#)
- installation, [Installing binutils](#)
- overview, [binutils](#)
- usage, [Using the GNU Assembler](#)

**GNU Binutils (see binutils)****GNU Compiler Collection, [GNU Compiler Collection \(GCC\)](#)**

- C support, [GNU C Compiler](#)
- C++ support, [GNU C++ Compiler](#)
- compatibility, [Language Compatibility](#), [ABI Compatibility](#), [Debugging Compatibility](#), [Other Compatibility](#), [C++11 Compatibility](#), [Compatibility Changes](#), [Fortran 2003 Compatibility](#), [Fortran 2008 Compatibility](#), [Fortran 77 Compatibility](#)
- documentation, [Additional Resources](#)

- features, [Main Features](#), [Status and Features](#), [New Features](#), [Fortran 2003 Features](#), [Fortran 2008 Features](#)
- Fortran support, [GNU Fortran Compiler](#)
- installation, [Installing the C Compiler](#), [Installing the C++ Compiler](#), [Installing the Fortran Compiler](#)
- overview, [GNU Compiler Collection \(GCC\)](#)
- usage, [Using the C Compiler](#), [Using the C++ Compiler](#), [Using the Fortran Compiler](#), [Preparing a Program for Debugging](#)
- version, [Main Features](#), [GNU Compiler Collection \(GCC\)](#), [Changes in GCC 4.7.0](#), [Changes in GCC 4.7.2](#)

### GNU Debugger, [GNU Debugger \(GDB\)](#)

- compatibility, [Compatibility Changes](#)
- documentation, [Additional Resources](#)
- features, [Main Features](#), [New Features](#)
- installation, [Installing the GNU Debugger](#)
- overview, [GNU Debugger \(GDB\)](#)
- preparation, [Preparing a Program for Debugging](#)
- usage, [Running the GNU Debugger](#), [Listing Source Code](#), [Setting Breakpoints](#), [Starting Execution](#), [Displaying Current Values](#), [Continuing Execution](#)
- version, [Main Features](#), [GNU Debugger \(GDB\)](#), [Changes in GDB 7.4](#), [Changes in GDB 7.5](#)

### GNU linker, [Using the GNU Linker](#)

- compatibility, [Compatibility Changes](#)
- documentation, [Additional Resources](#)
- features, [New Features](#), [GNU Linker](#)
- installation, [Installing binutils](#)
- overview, [binutils](#)
- usage, [Using the GNU Linker](#)

### gold (see GNU linker)

### gprof

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

### GSS (see Global Support Services)

## H

### Helgrind

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

### help

- accessing the Customer Portal, [Accessing the Customer Portal](#)
- getting help, [Do You Need Help?](#)
- Global Support Services, [Contacting Global Support Services](#)
- Red Hat Product Documentation, [Accessing Red Hat Product Documentation](#)

## L

**ld** (see GNU linker)

**linking** (see GNU linker)

## M

### Massif

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

### Memcheck

- overview, [Valgrind](#)
- usage, [Using Valgrind](#)

## N

### nm

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

## O

### objcopy

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

### objdump

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

### opannotate

- overview, [OProfile](#)
- usage, [Using OProfile](#)

### oparchive

- overview, [OProfile](#)
- usage, [Using OProfile](#)

### opcontrol

- overview, [OProfile](#)
- usage, [Using OProfile](#)

### opgprof

- overview, [OProfile](#)
- usage, [Using OProfile](#)

**ophelp**

- overview, [OProfile](#)
- usage, [Using OProfile](#)

**opimport**

- overview, [OProfile](#)
- usage, [Using OProfile](#)

**opjitconv**

- overview, [OProfile](#)
- usage, [Using OProfile](#)

**opreport**

- overview, [OProfile](#)
- usage, [Using OProfile](#)

**OProfile, [OProfile](#)**

- documentation, [Additional Resources](#)
- installation, [Installing OProfile](#)
- overview, [OProfile](#), [Changes in OProfile 0.154](#)
- usage, [Using OProfile](#)
- version, [Main Features](#), [OProfile](#)

**oprofiled**

- overview, [OProfile](#)
- usage, [Using OProfile](#)

**R****ranlib**

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

**readelf**

- features, [New Features](#)
- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

**Red Hat Customer Portal (see Customer Portal)****Red Hat Developer Toolset, [Red Hat Developer Toolset](#)**

- compatibility, [Compatibility](#)
- documentation, [Additional Resources](#), [Accessing Red Hat Product Documentation](#)
- features, [Main Features](#)
- installation, [Installing Red Hat Developer Toolset](#)
- overview, [About Red Hat Developer Toolset](#)
- subscription, [Getting Access to Red Hat Developer Toolset](#)

- support, [About Red Hat Developer Toolset](#)
- uninstallation, [Uninstalling Red Hat Developer Toolset](#)
- update, [Updating Red Hat Developer Toolset](#)

### Red Hat Enterprise Linux

- documentation, [Additional Resources, Accessing Red Hat Product Documentation](#)
- supported versions, [Compatibility](#)

### Red Hat Subscription Management

- subscription, [Using Red Hat Subscription Management](#)

### RHN Classic

- subscription, [Using RHN Classic](#)

## S

### scl (see Software Collections)

#### size

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

### Software Collections

- documentation, [Additional Resources, Accessing Red Hat Product Documentation](#)
- overview, [About Red Hat Developer Toolset](#)

#### stap

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

#### stap-merge

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

#### stap-prep

- overview, [SystemTap](#)
- usage, [Installing SystemTap](#)

#### stap-report

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

#### staprun

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

**stapsh**

- overview, [SystemTap](#)
- usage, [Using SystemTap](#)

**strings**

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

**strip**

- overview, [binutils](#)
- usage, [Using Other Binary Tools](#)

**support**

- Red Hat Developer Toolset, [About Red Hat Developer Toolset](#)

**SystemTap, [SystemTap](#)**

- documentation, [Additional Resources](#)
- installation, [Installing SystemTap](#)
- overview, [SystemTap](#)
- usage, [Using SystemTap](#)
- version, [Main Features](#), [SystemTap](#)

**V****Valgrind, [Valgrind](#)**

- documentation, [Additional Resources](#)
- installation, [Installing Valgrind](#)
- overview, [Valgrind](#)
- usage, [Using Valgrind](#)
- version, [Main Features](#), [Valgrind](#)